

Draft



Reference Model for Service Oriented Architectures

Working Draft 08, September 9, 2005

Document identifier:

wd-soa-rm-08

Location:

http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=soa-rm

Editors:

C. Matthew MacKenzie, Adobe Systems Incorporated, mattm@adobe.com
Ken Laskey, Mitre Corporation, klaskey@mitre.org
Francis McCabe, Fujitsu Limited, fgm@fla.fujitsu.com
Peter Brown, Individual, peter@justbrown.net
Rebekah Metz, Booz Allen Hamilton, metz_rebekah@bah.com

Abstract:

This ~~Service Oriented Architecture Reference Model~~ is an abstract framework for understanding significant entities and relationships ~~amongst~~ them within a service-oriented environment, and for the development of consistent standards or specifications supporting that environment. It is based on unifying concepts of SOA and may be used by architects developing specific services oriented architectures or ~~for education~~ and explaining SOA. A reference model is not directly tied to any standards, technologies or other concrete implementation details, but it does seek to provide a common semantics that can be used unambiguously across and between different implementations.

While service-orientation may be a popular concept found in a broad variety of applications, this reference model scopes itself to the field of software architecture.

Status:

This document is updated periodically on no particular schedule. Send comments to the editor(s).

Committee members should send comments on this specification to the soa-rm@lists.oasis-open.org list. Others should visit the SOA-RM TC home page at http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=soa-rm, and record comments using the web form available there.

Draft

36
37
38
39
40
41
42
43
44

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 SOA-RM TC web page at:

http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=soa-rm

The errata page for this specification is at:

http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=soa-rm.

Draft

44 **Table of Contents**

45 1 Introduction4
46 1.1 What is a reference model.....4
47 1.2 Audience4
48 1.3 How to use the reference model4
49 1.4 Notational Conventions5
50 1.5 Relationships to Other Standards5
51 2 The Reference Model.....7
52 2.1 Overview of model.....7
53 2.2 The Reference Model in Detail8
54 2.2.1 Service.....8
55 2.2.2 Service description9
56 2.2.3 Descriptions and Metadata11
57 2.3 Interacting with services.....12
58 2.3.1 Information model.....13
59 2.3.2 Behavior model14
60 2.3.3 Services in context15
61 2.4 Policies and Expectations16
62 2.4.1 Service Policy16
63 2.4.2 Services and expectations17
64 2.5 Service discoverability19
65 3 Conformance Guidelines20
66 4 References21
67 4.1 Normative21
68 4.2 Non-Normative21
69 Appendix A. Glossary22
70 Appendix B. Acknowledgments.....26
71 Appendix C. Notices.....27
72

1 Introduction

1.1 What is a reference model

A reference model consists of a set of clearly defined basic concepts, axioms and relationships within a particular problem domain, independently of specific implementations, conventions, activities or organizations. The purpose of a reference model is to facilitate the design of systems, to establish a common set of terminology as it applies to the domain and to encourage best practice where possible.

In the field of information technology, specific *architectures* may be developed to promote a common approach to solving particular problems. A group of such architectures – Service Oriented Architectures or SOAs – have developed over recent years with a specific mission to improve interoperability and collaboration between otherwise dissimilar services.

SOAs have received significant attention within the software design and development industry in recent times resulting in many conflicting definitions of “service-oriented architecture”. Whereas SOA architectural patterns (or *reference architectures*) may be developed to explain and underpin the generic design template supporting a specific SOA, a reference model is intended to provide an even higher level of commonality, with definitions that should apply to *any* SOA at all. The value of such a reference model is as a foundational work that can and should be used to develop architectural patterns and promote effective discourse on derived works.

The goal of this reference model document is to define the essence of the service oriented architecture paradigm, and emerge with a vocabulary and a common understanding of SOA. It should provide a normative reference that remains relevant for SOA as an abstract and powerful *model*, irrespective of the various and inevitable technology evolutions that we experience in this industry and that will impact on specific SOA *implementations*.

1.2 Audience

The intended audiences of this document non-exhaustively ~~include~~ include:

- Architects and developers designing, identifying or developing a system based on the service-oriented paradigm.
- Standards architects / analysts developing specifications that relate to or make use of the service-oriented paradigm.
- Chief Information Officers and other decision makers seeking a "consistent and common" understanding of service oriented architecture.

1.3 How to use the reference model

New readers are encouraged to read this reference model in its entirety. Concepts are presented in an order that the authors hope promote rapid understanding.

Draft

111

112 This section introduces the conventions, defines the audience and sets the stage for the rest of
113 the document. Non-technical readers are encouraged to read this information as it provides
114 background material necessary to understand the nature of reference models and their use.

115

116 Section 2 introduces the Reference Model for SOAs. First, the main axioms, key concepts and
117 relationships between those concepts are introduced followed by more detailed sections on the
118 main concepts: a *service* is defined along with *service description*. There then follows a section
119 detailing interaction between services, followed by service policies and expectations. Finally,
120 being the key to a service's actual use, the concept of service discoverability is introduced.

121

122 This section is provided for the benefit of multiple audiences:

- 123 • Non-technical readers may use this section to gain an explicit understanding of the core
124 principles of SOA.
- 125 • Architects are encouraged to use this section as guidance for developing specific service
126 oriented architectures. Section 2 and its subsections are designed to provide guidance
127 for consistent logical divisions of components within architectures. It also helps architects
128 adhere to the basic principles of service-oriented design.

129

130 Section 3 addresses what it might mean for an SOA-based system to be conformant with this
131 reference model.

132

133 The glossary provides definitions of terms which are relied upon within the reference model
134 specification but do not necessarily form part of the specification itself.

135

136 1.4 Notational Conventions

137 The key words *must*, *must not*, *required*, *shall*, *shall not*, *should*, *should not*, *recommended*, *may*,
138 and *optional* in this document are to be interpreted as described in [RFC2119].

139

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

141 1.5 Relationships to Other Standards

142 Due to its nature, this reference model may have an implied relationship with any group that:

- 143 • Considers its' work "Service Oriented"; and/or
- 144 • Makes (publicly) an adoption statement to use this SOA Reference Model of this TC as a
145 base or inspiration for their work when complete.

146

147 Additionally, there are a large number of standards and technologies that are related by the fact
148 they claim to be or are "service oriented".

Draft

149 Any work that aligns with the functional areas of SOA such as the service, service description,
150 advertising mechanism, service data model or service contract are likely to be directly related.

151

152 The reference model does not endorse any particular service-oriented architecture, or attest to
153 the validity of third party reference model conformance claims.

154

155 2 The Reference Model

156 The reference model for service-oriented architectures describes concepts and relationships that
157 are fundamental in describing SOA architecture patterns (i.e. SOA reference architectures) and
158 specific SOA architectures applied to the solution of specific problems. In general, a service-
159 oriented architecture represents a uniform means to discover and access distributed services that
160 invoke functionality ~~which~~ produces desired effects with measurable preconditions and
161 expectations. ~~The~~ services hide implementation details but have associated service descriptions
162 to provide sufficient information to understand the technical and business requirements for
163 invoking the service. The actual decision (or agreement) to invoke a service ~~often is~~ contingent
164 on understanding and complying with those requirements.

165

166 While such a description of SOA gives a flavor for why it is of interest, it is not sufficient for
167 understanding the primary SOA concepts that must be utilized in designing a SOA and effectively
168 using an SOA. The remainder of this section introduces the main concepts and a detailed
169 discussion of the concepts and their relationships are in the sections that follow.

170 2.1 Overview of model

171 A key concept of SOA is that of a **service**. In general, people and organizations create
172 capabilities to solve or support the solution of problems they face in the course of their business.
173 SOA is conceived as a way of making those capabilities visible and supporting standard means of
174 access so the existing capabilities can be reused or new capabilities can be readily substituted to
175 improve the solutions. A service is a means to access such capabilities.

176

177 To use a service, it is necessary to know it exists, what is accomplished if the service is invoked,
178 how ~~to invoke~~ the service, and other characteristics to allow a prospective consumer to decide if
179 the service is suitable for the current needs and if ~~the~~ consumer satisfies any requirements of the
180 service provider to be permitted access. Such information constitutes the **service description**.

181

182 Services are accessed in order to achieve particular effects. However, the nature of SOAs are
183 that there is an arm's length relationship between service providers and consumers. As a result,
184 there is a distinction to be drawn between the public interactions with a service and the private
185 actions of the service provider and consumer. ~~An important reason for~~ the scalability and security
186 attributes of SOAs is that the distinction ~~promotes~~ independence ~~between~~ service participants.
187 We can focus on the public aspects of using a service by examining the **conditions** of using a
188 service and the **expectations** that arise as a result of using the service. We loosely associate the
189 service conditions with the **service policies** and the expectations with **service contracts**.

190

191 Another key concept in SOA is that of **service interaction**. Although services are accessed in
192 order to achieve particular desired effects, this is effected by exchanging information between
193 service providers and consumers. Typically this is by exchanging messages using a standardized
194 protocol; however, there are many modalities possible for using services

Draft

195

196 Finally we identify **discoverability** as a key concept of SOAs. Discoverability refers to the
197 possibility and mechanisms by which service consumers and providers can be brought together.
198 There are many possible mechanisms by which discoverability may be achieved; ~~SOAs are not~~
199 ~~limited to~~ registries or repositories of service descriptions ~~although these are~~ undoubtedly
200 powerful means of achieving it. Discoverability ~~itself is~~ a key concept for SOAs.

201 2.2 The Reference Model in Detail

202 2.2.1 Service

203 A service is a mechanism to enable access to a set of capabilities, where the access is provided
204 using a prescribed interface and is exercised consistent with constraints and policies as specified
205 by the service description. A service is provided by one entity for use by others, but the eventual
206 consumers of the service may not be known to the service provider and may demonstrate uses of
207 the service beyond the scope originally conceived by the provider.

208

209 A service is invoked through a service interface, where the interface comprises the specifics of
210 how to access the underlying capabilities. There are no constraints on what constitutes the
211 underlying capability or how access is implemented by the service provider. Thus, the service
212 could carry out its described functionality through one or more automated and/or manual
213 processes that themselves could invoke other available services. A service is unique in that its
214 implementation is hidden from the service consumer except for (1) the **information** model
215 exposed through the published service interface and (2) any information included as metadata to
216 describe aspects of the service which are needed by service consumers to determine whether a
217 given service is appropriate for the consumer's needs. The consequence of ~~exercising~~
218 is one or more real world effects. The effects may include

219

220 (1) information returned in response to a request for that information, including information
221 returned as a result of prior interactions with the service,

222 (2) processing done in response to a request to change the state of **identified** entities, or

223 (3) some combination of (1) and (2).

224

225 Note, the user in (1) does not typically know how the information is generated, e.g. whether it is
226 extracted from a database or generated dynamically; in (2), the user does not typically know how
227 the state change is effected. In either case, the service consumer would need to provide input
228 parameters defined (either required or optional) by the service and the service would return
229 information, status indicators, or error descriptions, where both the input and output are as

¹² http://en.wikipedia.org/wiki/Turing_completeness

Draft

230 described by the **information** model exposed through the published service interface. Note that
231 the service may be invoked without requiring information input from the consumer (other than a
232 command to initiate action) and may accomplish its functions without providing any return or
233 feedback to the consumer.

234 The **description** of the service concept has emphasized a distinction between a capability that
235 represents some functionality created to address a problem or a need and the service that forms
236 the point of access to bring that capability to bear in the context of SOA. It is assumed the
237 capability was created and exists outside of SOA and one of the major benefits of SOA is
238 enabling the capability to be applied to an expanded realm of relevant problems. In actual use,
239 maintaining this distinction may not be critical (i.e. the service may be talked about in terms of
240 being the capability) but the separation is pertinent in terms of a clear expression of the nature of
241 SOA and the value it provides.

242 2.2.2 Service description

243 The service description represents the information needed to use a service. It may be considered
244 part of or the complete set of the metadata associated with a service (see Section 2.2.3) but in
245 any case, the service description overlaps and shares many common properties with service
246 metadata. In most cases, there is no one "right" set of metadata but rather the metadata content
247 depends on the context and the needs of the parties using the associated entity. The same holds
248 for a service description. While there are certain elements that are likely to be part of any service
249 description, most notably the **information** model, many elements such as function and policy may
250 vary. However, ~~the mechanisms to specify the~~ service description should be represented through
251 use of a standard, reference-able format **that accommodate the necessary variations** and
252 lend themselves to common processing tools (such as discovery engines) to make use of the
253 service description.

254

255 While the concept of a SOA supports use of a service without the service consumer needing to
256 know the details of the service implementation, the service description makes available critical
257 information a consumer needs to decide to use a service and ~~then to affect that use~~. In
258 particular, a service consumer must ~~know~~:

259

- 260 1. The service exists and is available;
- 261 2. The service performs a certain function or set of functions consistent with technical
262 assumptions that underlie its functions;
- 263 3. The service operates under a specified set of constraints and policies;
- 264 4. The service will (to some extent) comply with policies as prescribed by the service
265 consumer;
- 266 5. How to interact with the service in order to achieve the required objectives, including the
267 format and content of information exchanged between the service and the consumer and
268 the sequences of information exchange that may be expected.

269

270 Subsequent sections of this document will deal with these aspects of a service in details but the
271 following subsections will describe their relation to the service description.

Draft

272 2.2.2.1 Service Availability

273 Item 1 refers to the key requirement for a service description to include sufficient data to permit a
274 service consumer and service provider to **physical** exchange information. This might **range from**
275 metadata such as the location of the service and what information protocols it supports and
276 requires **to** whether the service is currently available or not.

277 2.2.2.2 Service Functionality

278 Item 2 relates to the need to unambiguously express the function(s) of the service and the real
279 world effects (see section 2.4) that result from it being invoked. This portion of the description
280 needs to be expressed in a way that is generally understandable by service consumers but able
281 to accommodate a vocabulary that is sufficiently expressive for the domain for which the service
282 provides its functionality. The description of functionality may include, among other possibilities,
283 a textual description intended for human consumption or identifiers or keywords referenced to
284 specific machine-process-able definitions. For a full description, it may be useful to indicate
285 multiple identifiers or keywords from a number of different collections of definitions.

286

287 Part of the description of functionality may include underlying technical assumptions that
288 determine the limits of functionality exposed by the service or of the underlying capability. For
289 example, the amounts dispensed by an automated teller machine (ATM) are consistent with the
290 assumption that the user is an individual rather than a business. To use the ATM, the user must
291 not only adhere to the policies and satisfy the constraints of the associated financial institution
292 (see section **1.1** for how this relates to service description and section 2.4 for a detailed
293 discussion) but the user is limited to withdrawing certain fixed amounts of cash and a certain
294 number of transactions in a specified period of time. The financial institution, as the underlying
295 capability, does not have these limits but the service interface **it** exposes to its customers does,
296 consistent with its assumption of the needs of the intended user. If the assumption is not valid,
297 the user may need to use another service to access the capability.

298

299 2.2.2.3 Policies Related to a Service

300 Items 3 and 4 relate to the service description's support for associating constraints and policies
301 with a service and providing necessary information for prospective consumers to evaluate if a
302 service will act in a manner consistent with the consumer's constraints and policies.

303

304 In some situations the consumer may similarly provide an indication of its constraints and policies
305 to support a service's need to do a similar evaluation of suitability. Thus, both prospective
306 consumers and providers are likely to use the service description (and the consumer description)
307 to mutually establish what section 2.3.3 refers to as the *execution context*.

308

309 2.2.2.4 Service Interface

310 The service interface is the means referred to in Item 5 for interacting with a service. It includes
311 the specific protocols, commands, and information exchange by which actions are initiated that

Draft

312 result in the real world effects as specified through the service functionality portion of the service
313 description.

314

315 The specifics of the interface are syntactically represented in a standard reference-able format.
316 These prescribe what information needs to be provided to the service in order to exercise its
317 functionality and/or the results of the service invocation to be returned to the service requester.
318 This logical expression of the set of information items associated with the consumption of the
319 service is often referred to as the service's **information** model. ~~Note, specifying~~ the particulars of
320 the standard reference-able format is beyond the scope of the reference model ~~but~~ requiring that
321 mechanisms be available to define and retrieve such definitions ~~are~~ fundamental to the SOA
322 concept. ~~Also note that the service may be invoked without requiring information input from the~~
323 ~~consumer (other than a command to initiate action) and may accomplish its functions without~~
324 ~~providing any return or feedback to the consumer.~~

325 While this discussion refers to a standard reference-able syntax for service descriptions, we do
326 not specify how the consumer accesses the interface definition nor how the service itself is
327 accessed. However, it is assumed that for a service to be usable, ~~its interface must be~~
328 ~~represented in a format that allows interpretation of the interface information.~~

329 2.2.2.5 An Example of Using Information Contained in the Service 330 Description

331 The following example may help to clarify the concepts related to service and service description.
332 ~~To access electricity generated by the local electric utility, the service interface is the wall outlet~~
333 ~~and to use the service I need to understand what kind of plug fits the outlet. The utility assumes I~~
334 ~~will plug in devices that are compatible with the voltage they are providing and my assumption is I~~
335 ~~can safely plug in devices without these being damaged. If I am a home or business user, a~~
336 ~~constraint is I must establish an account and the contract I have with the electric utility is they will~~
337 ~~meter my usage and I will pay at a rate they prescribe. If I am a visitor to someone with a~~
338 ~~contract, the utility does not have a contract with me (and I do not have to satisfy the initial~~
339 ~~account constraint) but I still must be compatible with the service interface. The utility policy may~~
340 ~~be that in the event of high use by the community, the utility may reduce voltage or institute rolling~~
341 ~~blackouts. My implied policy is I may complain to my legislative representative if this happens~~
342 ~~frequently. The resource is the utility's ability to generate and distribute electricity, and the~~
343 ~~service is my getting access to that electricity. The resource would exist if every device was~~
344 ~~required to be hardwired to the electric utility's equipment but this would result in a very different~~
345 ~~service with a very different interface.~~

346 2.2.3 Descriptions and Metadata

347 One of the hallmarks of a Service Oriented Architecture is the degree of documentation and
348 description associated with it; particularly *machine processable descriptions* – otherwise known
349 as *metadata*.

350

351 The purpose of this metadata is to facilitate integration, particularly across ownership domains.
352 By providing public descriptions, it makes it possible for potential participants to construct
353 applications that use services and even offer compatible services with minimal human-level
354 contact between them.

Draft

355 2.2.3.1 The roles of description

356 An important additional benefit of metadata – as opposed to informal natural language
357 descriptions – is that it potentially permits automation in the creation of software. Both service
358 providers and service consumers can benefit from such automation – reducing the cost of
359 developing such systems.

360

361 For example, metadata can be used as a basis of discovery in dynamic systems, it can assist in
362 managing a service, validating and auditing uses of services may also be simplified by rich
363 metadata. It can help to ensure that requirements and expectations regarding the content of any
364 data interchanged are properly interpreted and fulfilled.

365 2.2.3.2 The limits of description

366 There are well-known theoretic limits on the effectiveness of descriptions – it is simply not
367 possible to completely specify in a completely unambiguous manner the precise semantics of a
368 service. (This is Gödel's incompleteness result in another guise)

369

370 Another way of stating the above is that there will always be unstated assumptions made by the
371 describer of a service that must be implicitly shared by readers of the description. This applies to
372 machine processable metadata as well as to human readable documentation.

373

374 Luckily, such precision is not normally necessary, either – what is required is sufficient precision to
375 enable required functionality.

376

377 Another kind of limit of descriptions is more straightforward: describing a service (for example)
378 does not eliminate the requirement for making a choice. For example, a service directory might
379 have the descriptions of many services – provided by many organizations. An automatic search
380 of that directory is therefore likely to return multiple responses to any mechanical search criteria.
381 At some point this set of responses has to be converted into a choice of a single service in order
382 for a service consumer (say) to perform its function. In a multi-vendor scenario, that choice must
383 also take into account real world aspects of the service – such as who the provider of the service
384 is and whether the service consumer can or should trust the provider. It is unlikely that such
385 factors can be easily and securely encoded in descriptions and search criteria.

386 2.3 Interacting with services

387 Interacting with a service involves exchanging information with the service and performing actions
388 against the service. In many cases, this is accomplished by sending and receiving messages to
389 and from the service end-point; but there are other modes possible that do not involve explicit
390 message sending. However, for simplicity, we often refer to message exchange as the primary
391 mode of interaction with a service. Together the forms of information exchanged and the kinds of
392 actions performed form the **service interface** – see section 2.2.2.

393

394 There are three key concepts that are important in understanding what it is involved in interacting
395 with services – the **information model**, the **process model** and the **execution context**.

Draft

396 2.3.1 Information model

397 The information model of a service is a characterization of the information that may be exchanged
398 with the service.

399

400 The scope of the information model includes the format of documents and messages, the
401 structural relationships within those documents and also the ontologies of terms used within those
402 documents.

403

404 ~~We do not, however, generally include within the information model of a service~~ the information
405 and data that might be stored or internally manipulation by a service. That is part of the service
406 implementation.

407 2.3.1.1 Structure

408 Knowing the representation, structure and form of information required is a key initial step in
409 ensuring effective interactions with a service. There are several levels of such structural
410 information; ranging from the encoding of character data, through the use of formats such as
411 XML, SOAP and schema-based representations.

412 2.3.1.2 Ontology

413 Particularly for messages, an important aspect of the service information model is the
414 interpretation of strings and other tokens in the data. Loosely, one might partition the
415 interpretation of a message into structure (syntax) and ontology (semantics); although both are
416 part of the information model.

417

418 A described information model typically has a great deal to say about the form of messages,
419 about the types of the various components of messages and so on. However, pure type
420 information is not sufficient to completely describe the appropriate interpretation of data. For
421 example, within an address structure, the city name and the street name are typically given the
422 same type – some variant of the string type. However, city names and street names are not really
423 the same type of thing at all. Distinguishing the correct interpretation of a city name string and a
424 street name string is not possible using type-based techniques – it requires additional information
425 that cannot be expressed purely in terms of the structure of data.

426

427 ~~Ontologies are formal descriptions of sets of terms in terms of the~~ relationships between them.
428 Most commonly, the relationships are class relationships – one term represents a concept that is
429 a sub-class of another. However, relationships are not limited to the sub-class relationships; other
430 aspects of concepts can also be usefully represented; such as the range of possible values given
431 property can take and whether the property is functional or not.

432

433 The role of explicit ontologies is to provide a firm basis for selecting correct interpretations for
434 tokens in messages. For example, in the address example above, an ontology can be used to
435 capture the appropriate distinction between street name and city name; so much so that in many

Draft

436 cases it is possible to automatically map the contained information from one representation to
437 another.

438

439 More specifically, for a service to be consistent, the service should make consistent use of terms
440 as defined in one or more ontologies. Of course, specific domain semantics are beyond the scope
441 of this reference model; however, the reference model does address the requirement that the
442 service interface enable providers and consumers to unambiguously identify relevant definitions
443 for their respective domains.

444 2.3.2 Behavior model

445 The second key requirement for successful interactions with services is knowledge of the
446 behavioral or process aspects of the service. Loosely, this can be characterized as knowledge of
447 the actions on, responses to and temporal dependencies between actions on the service.

448

449 For example, in a News subscription service, a successful use of the service involves initially
450 registering a subscription with the service; which will then be followed by an irregular series of
451 one-way news items. Key to using the service is the knowledge that you must first register your
452 preferences and then you will get messages without further prompting.

453

454 Another example is a service that supports updating a balance with a transaction. Such services
455 are typically *idempotent*: i.e., they will not change their state should a subsequent interaction be
456 attempted for the same transaction. The behavioral model of the account update service then
457 consists of an initial communication – incorporating the transaction to log – followed by a
458 response which includes the new balance.

459

460 2.3.2.1 Process Model

461 It is fairly common to partition the process model associated with a service into two levels: the
462 particular sequences of operations needed to achieve single service exchanges and longer term
463 transactions. These two levels may be nested – a long running transaction is often composed of
464 sequences of exchange patterns.

465

466 For example, in a publish-and-subscribe service, there are individual operations dealing with
467 registering a new subscription (say) and publishing a new notice (say). The longer view of a given
468 service considers the total sequence of notifications associated with a given subscription.
469 Another concept that may be featured in a process model is the transactional structure of a
470 service (c.f. ACID analysis of processes).

471

472 Note that although the existence of a process model is fundamental to this Reference Model, its
473 extent is not defined. In some architectures the process model will include aspects that are not
474 strictly part of this reference model – for example we do not address the orchestration of multiple
475 services – although orchestration and choreography may be part of the process model of a given
476 architecture. At a minimum, the process model must cover the interactions with the service itself.

Draft

477

478 Choosing an appropriate representation of process models is a fine art; a representation system
479 that can express sequences and dependencies is often Turing complete² – i.e., is effectively a
480 programming language. The problem with Turing complete representations of processes is that
481 processing such descriptions quickly becomes intractable for non-trivial process models. For
482 example, the task of comparing two processes is a difficult exercise that is provably impossible in
483 the general case. On the other hand, without some such expressive power it can be difficult to
484 capture the required dependencies that are a natural part of process descriptions.

485

486 However, showing that two process models are equivalent is not the only requirement for
487 representing process models. A more common requirement is simply to be able to identify the
488 appropriate steps that must be followed for a successful interaction. This is analogous to following
489 a recipe or executing a program – a task that is easily mechanizable.

490

491 2.3.2.2 Behavior

492 The **behavioral** model of a service is about the behavior that results in interactions with the
493 service. Of course, a great portion of the behavior of a service may be private; however, the
494 expected public view of a service surely includes the implied behavior of the service.

495

496 For example, in a service that represents a bank account, it is not sufficient to know that to use
497 the service you need to exchange a given message (with appropriate authentication tokens). It is
498 also of the essence that using the service may actually affect the bank account – withdrawing
499 cash from it for example.

500

501 The behavior of a service is closely connected to its intended real-world effect; although not
502 identical to it. In general, we can state that the behavior of a service (an attempt to withdraw cash
503 from an account) results in an intended (or occasionally unintended) effect in the world: the
504 account's balance is lower.

505

506 2.3.3 Services in context

507 In an implementation, services are associated with an **execution context**. Or, another way of
508 expressing this is to consider that there is a distinction between a potential service and an actual
509 service that is capable of being interacted with. An actualized service has an execution context
510 that determines many of the properties of the service; including attributes such as security.

511

512 For example, suppose that it were important that a given service was always executed in an
513 authenticated context – i.e., that the service provider and the service consumer have
514 authenticated themselves to each other. The details of how authentication is performed are not
515 our concern here. That authentication context is an example of a particular execution context that
516 applies to the service.

517

Draft

518 The execution context is a touchstone for many aspects of the service – what policies are in force
519 for example, whether it is available, and so on.

520 2.4 Policies and Expectations

521 In the absence of intimate knowledge of the implementation of service providers and consumers
522 a way of characterizing the use of a service is via the concepts of **policies** and **expectations**. We
523 can understand what it means to interact with a service by examining the conditions on the use of
524 the service and the expected results of using it.

525

526 Broadly speaking, a policy represents some form of constraint or condition on the use,
527 deployment or description of an owned entity. We are focused primarily on the concept of policy
528 as it applies to services.

529

530 On the other hand, the expectations associated with a service revolve around the consequences
531 of interacting with the service. Normally, there is an expected **real world effect** as a result of
532 using a service – such as depositing money into an account. However, it can be difficult to
533 characterize the real world effect of using a service, since the actions performed by a service
534 implementation are inherently private to that provider.³ Instead it may be more effective to
535 consider the expectations for future interactions with services.

536 2.4.1 Service Policy

537 Abstractly, a policy is a statement of the obligations, constraints or other conditions of use of a
538 given service that expresses intent on the part of a participant. More particularly, policies are a
539 way for expressing the relationship between the **execution context** and the **information** and
540 **process models** associated with the service.

541

542 Conceptually, there are three aspects of policies: the policy assertion, the policy owner
543 (sometimes referred to as the policy subject) and policy enforcement.

544

545 For example, the assertion: “All messages are triple-DES encrypted” is an assertion constraining
546 the forms of messages. As an assertion, it is measurable: it may be true or false depending on
547 whether the traffic is actually encrypted or not. Note that policy assertions are often about the way
548 the service is realized; i.e., they are about the relationship between the service and its execution
549 context.

550

³ A similar analysis applies to service consumers: just how a consumer of a service decides which requests to make is something that the service provider cannot determine.

Draft

551 A policy always represents a participant's point of view. An assertion becomes the policy of a
552 participant when they make it their policy – this linking is normally not part of the assertion itself.
553 For example, if the service consumer declares that "All messages are triple-DES encrypted", then
554 that reflects the policy of the service consumer. This policy is one that may be asserted by the
555 service consumer independently of any agreement from the service provider.

556

557 Finally, a policy may be enforced. Techniques for the enforcement of policies depend on the
558 nature of the policy. From a conceptual point of view, service policy enforcement amounts to
559 ensuring that the assertion is consistent with the real world. An unenforceable policy is not really
560 a policy; it would be better described as a wish.

561

562 Policies potentially apply to many aspects of SOA: security, privacy, manageability, Quality of
563 Service and so on. Beyond such infrastructure-oriented policies, participants may also express
564 business-oriented policies – such as hours of business, return policies and so on.

565

566 Policy assertions may be, but need not be, written down in a formal machine processable form.
567 The importance of such a machine processable form of policy depends on the purpose and
568 applicability of the policy. In particular, where a policy declaration might affect whether a particular
569 service is used or not, then such policies should be expressed in machine-processable form.

570

571 Languages that permit policy assertions also range in expressivity from simple propositional
572 assertions to modal logic rules. However, the Reference Model is neutral to how a policy is
573 represented.

574

575 A natural point of contact between service participants and policies associated with the service is
576 in the service description. It would be natural for the service description to contain references to
577 the policies associated with the service.

578 2.4.2 Services and expectations

579 There is nearly always a particular purpose associated with interacting with a service – the
580 service consumer is trying to achieve some result by invoking the service, as is the service
581 provider. At first sight, such a goal can often be expressed as "trying to get the service to do
582 something" – this is sometimes known as the **real world effect** of using a service. For example,
583 an airline reservation service can be used in order to book seats on a flight.

584

585 However, inherent to the concept of SOA is an arm's length approach to the relationship between
586 service providers and consumers where there are minimum assumptions made by consumers
587 about how a service is provided, and conversely minimum assumptions made by service
588 providers about the connectivity of consumers. This separation is key to achieving large-scale
589 systems and also to managing the evolution of such systems.

590

591 In keeping with this assumption, a more effective way of capturing the purpose of using a service
592 is via the concept of **expectations**. I.e., rather than trying to ensure that the airline reservation

Draft

593 service has recorded our booking, we are more interested in knowing that when we arrive at the
594 airport, the airline will agree that we do indeed have a seat on the flight.

595

596 Expectations revolve around communication and future interactions much more than present time
597 actions. Of course, in order for the airline to know that the seat is confirmed it will likely use some
598 kind of system for recording the reservation; but, by minimizing assumptions about how the airline
599 fulfils its contracts, we maximize the potential for smooth interoperation.

600

601 One way to characterize the expectations associated with a service interaction involves the
602 message traffic exchanged with the service. In a manner that is completely analogous to the
603 service interface, we can define expectations in terms of the kinds of information that will be
604 provided subsequently by a service – as opposed to the information that is required for a current
605 interaction. For example, a successful interaction with a courier or package delivery service might
606 result in a tracking number. The expectation is that presenting that tracking number to the
607 appropriate service will result in information about the current whereabouts of the package being
608 delivered.

609

610 The expectations arising from a use of a service may be described in much the same ways that
611 policies are described, except that the natural container for this is the **service contract**.

612 **2.4.2.1 Service Contract**

613 Where a policy is associated with the point of view of individual participants, a contract represents
614 an agreement between two or more participants. Like policies, contracts can cover a wide range
615 of aspects of services: quality of service agreements, interface and choreography agreements
616 and commercial agreements.

617

618 Thus, following the analysis above, a service contract is a measurable assertion that governs the
619 requirements and expectations of two or more parties. Unlike policy enforcement, which is
620 usually the responsibility of the policy owner, contract enforcement may involve disputes between
621 the parties to the contract. The resolution of such disputes may involve appeals to higher
622 authorities.

623

624 Like policies, contracts may be, but need not be, expressed in a machine processable form.
625 Where a contract is used to codify the results of a service interaction, it is good practice to
626 represent it in a machine processable form – that would facilitate automatic service composition
627 for example. Where a contract is used to describe over-arching agreements between service
628 providers and consumers then the priority is likely to make such contracts readable by people.

629

630 A variant of the policy concept is the agreement or **contract**. A contract has all the same features
631 as a policy with one key addition: the concept of agreement – contracts are policies that have
632 been agreed to by participants governed by the policy; policies do not need agreement only
633 enforcement.

Draft

634 **2.5 Service discoverability**

635 A key concept of the SOA Reference Model is the discoverability of services; which is an
636 important aspect of bringing together the service provider and consumer: A service provider must
637 be capable of making details of the service (notably service description and policies) available to
638 potential customers; and customers must be capable of finding that information.

639

640 This might (and commonly does) involve a service provider entering the service description into a
641 service registry and the service consumer searching for an appropriate match to their needs. In a
642 pure P2P architecture there would be no registry at all. The SOA concept of discoverability is not
643 restricted to any single mechanism.

644

645 Service Discoverability requires that the service description and policy – or at least a suitable
646 subset thereof – be available in such a manner and form that, directly or indirectly, an awareness
647 of the existence and capabilities of the service can become known to potential consumers. The
648 extent to which the discovery is “pushed” by the service provider, “pulled” by a potential
649 consumer, subject to a probe or another method, will depend on many factors.

650

651 For example: a service provider may advertise and promote their service by either including it in a
652 service directory or broadcasting it to all consumers; potential consumers may broadcast their
653 particular service needs in the hope that a suitable service responds with a proposal or offer or a
654 service consumer might also “probe()” an entire network to determine if suitable services exist.
655 When the demand for a service is higher than the supply, then by advertising their needs,
656 potential consumers are likely to be more effective than service providers advertising offered
657 services.

658

659 One way or another, the potential consumer must acquire a sufficient description as a prelude to
660 evaluating whether the service matches their expectations and, if so, how to proceed to establish
661 a contract and invoke the service.

662

663 In some contexts there are advantages to a registry model in which service capabilities are
664 advertised, and then matched against requests of service. Such models should ideally allow for
665 the capture of information both of service offers and service requests, so that it can be used in a
666 reverse registry that records needs and queries on offers as well as, or instead of, records of
667 service offers.

668

669 Specific SOA reference architectures and implementations will prescribe the mechanisms for
670 actual service discovery, ensuring a service’s presence and availability, and failure conditions and
671 error handling.

Draft

672 3 Conformance Guidelines

673 The authors of this reference model envision that architects may wish to declare their architecture
674 is conformant with this reference model. Conforming to a Reference Model is not generally an
675 easily automatable task – given that the Reference Model’s role is primarily to define concepts
676 that are important to SOAs rather than to give guidelines for implementing systems.

677 However, we do expect that any given Service Oriented Architecture will reference the concepts
678 outlined in this specification. As such, we expect that any design for a system that adopts the
679 SOA approach will

- 680 • Have entities that can be identified as services as defined by this Reference Model,
- 681 • Such entities will have descriptions associated with them,
- 682 • Service entities will have identifiable interaction models, including models of the
683 information exchanged by the services and the temporal behavior of the services
- 684 • It should be possible to identify a means by which consumers of services and providers
685 of services are able to engage; and
- 686 • That there will be identifiable aspects of service entities that correspond to the policies
687 relating to the conditions of use of services and to the expectations that result from
688 interacting with services.

689 It is not appropriate for this specification to identify *best practice* with respect to building SOA-
690 based systems. However, the ease with which the above elements can be identified within a
691 given SOA-based system could have significant impact on the scalability, maintainability and
692 ease of use of the system.

Draft

693 **4 References**

694 **4.1 Normative**

695 **[RFC2119]** S. Bradner, *Key words for use in RFCs to Indicate Requirement Levels*,
696 <http://www.ietf.org/rfc/rfc2119.txt>, IETF RFC 2119, March 1997.
697

698 **4.2 Non-Normative**

699 **[W3C WSA]** W3C Working Group Note "Web Services Architecture",
700 <http://www.w3.org/TR/ws-arch/> , 11 February 2004

Draft

701 Appendix A. Glossary

702 Terms that are used within this Reference Model are often also found in other specifications. In
703 order to avoid potential ambiguity, this glossary locally scopes the definitions of those terms for
704 the purpose of this Reference Model and thus overrides any other definitions.

705

706 Advertising (or Announcement of Availability)

707 A means of conveying the existence of and sharing awareness about a service to potential
708 consumers.

709

710 Agent (requester or provider)

711 An entity acting on behalf and with the authority of another entity and charged to fulfill a task.

712

713 Architecture

714 A set of artifacts (that is: principles, guidelines, policies, models, standards and processes) and
715 the relationships between these artifacts, that guide the selection, creation, and implementation of
716 solutions aligned with business goals.

717 Software architecture is the structure or structures of an information system consisting of entities
718 and their externally visible properties, and the relationships among them.

719

720 Authentication

721 The act by which an agent establishes – to an agreed level of confidence – the identity of another
722 entity.

723

724 (Service) Consumer

725 An entity which intends to make use of a service.

726

727 Contract

728 The syntactic, semantic and logical constraints governing the use of a service.

729

730 Data Model

731 A Data Model is the abstract paradigm used in the invocation and consumption of a service. It is
732 expressed as a set of information items associated with the use of a service.

733

734 Discovery

735 The act of detecting and gaining understanding of the nature of a service.

736

Draft

- 737 Encapsulation
738 The act of hiding internal specifications of an entity from the user of that entity, in such a way that
739 the internal data and methods of the entity can be changed without changing the manner in which
740 the entity is used. What is seen by the user is only an interface, or service.
741
- 742 Framework
743 A set of assumptions, concepts, values, and practices that constitutes a way of viewing the
744 current environment.
745
- 746 Interface
747 A named set of operations that characterize the behavior of an entity.
748
- 749 Mediation
750 The transformation, routing, validation and processing of messages.
751
- 752 Message
753 A serialized set of data that is used to convey a request or response from one party to another.
754
- 755 Metadata
756 A set of properties of a given entity which are intended to describe and/or indicate the nature and
757 purpose of the entity and/or its relationship with others.
758
- 759 Negotiation
760 A process that seeks to establish an acceptable basis for a contract between agents for the
761 provision of a service.
762
- 763 Ontology
764 Represents an agreement within a specific environment of the meanings to be associated with
765 different concepts and their relations to each other.
766
- 767 Opaqueness
768 The extent to which an agent is able to interact successfully with a service without detecting how
769 the service is implemented.
770
- 771 Policy
772 A statement of obligations, constraints or other conditions of use of a given service. When a
773 specific set of entities accept such a policy, a contract is usually established.
774

Draft

775 Reference Model

776 A reference model is an abstract framework for understanding significant relationships among the
777 entities of some environment that enables the development of specific architectures using
778 consistent standards or specifications supporting that environment.

779 A reference model is based on a small number of unifying concepts. A reference model is not
780 directly tied to any standards, technologies or other concrete implementation details, but it does
781 seek to provide a common semantics that can be used unambiguously across and between
782 different implementations.

783

784 (Service) Requester or provider

785 An agent that interacts with a service in order to achieve a goal

786

787 Security

788 A set of policies and measures designed to ensure that agents in an environment can only
789 perform actions that have been allowed. Security in a specific environment is an agreed
790 compromise between meeting the needs of agents and maintaining the integrity of the
791 environment.

792

793 Semantics

794 A conceptualization of the implied meaning of information, shared between the service consumer
795 and the service provider, that requires words and/or symbols within a usage context.

796

797 Service

798 A behavior or set of behaviors offered by one entity for use by another according to a policy and
799 in line with a service description.

800

801 Service description

802 A set of information describing a service, sufficient to allow a potential consumer to ascertain,
803 where appropriate:

804 - the identity of (and/or information about) the service provider;

805 - the policies, parameters and terms of use of the service;

806 - the procedures and constraints governing invocation of the service,

807 and thus determine whether the service meets the expectations and requirements of the
808 consumer. Acceptance of the service description by a consumer does not of itself imply a contract
809 to use the service.

810

811 Service Oriented Architecture (SOA)

812 A software architecture of services, policies, practices and frameworks in which components can
813 be reused and repurposed rapidly in order to achieve shared and new functionality. This enables

Draft

814 rapid and economical implementation in response to new requirements thus ensuring that
815 services respond to perceived user needs.
816 SOA uses the object-oriented principle of encapsulation in which entities are accessible only
817 through interfaces and where those entities are connected by well-defined interface agreements
818 or contracts.
819
820
821
822

Draft

823 **Appendix B. Acknowledgments**

824 The following individuals were members of the committee during the development of this
825 specification:

826 [TODO: insert cte. Members]

827

Draft

828

Appendix C. Notices

829 OASIS takes no position regarding the validity or scope of any intellectual property or other rights
830 that might be claimed to pertain to the implementation or use of the technology described in this
831 document or the extent to which any license under such rights might or might not be available;
832 neither does it represent that it has made any effort to identify any such rights. Information on
833 OASIS's procedures with respect to rights in OASIS specifications can be found at the OASIS
834 website. Copies of claims of rights made available for publication and any assurances of licenses
835 to be made available, or the result of an attempt made to obtain a general license or permission
836 for the use of such proprietary rights by implementers or users of this specification, can be
837 obtained from the OASIS Executive Director.

838 OASIS invites any interested party to bring to its attention any copyrights, patents or patent
839 applications, or other proprietary rights, which may cover technology that may be required to
840 implement this specification. Please address the information to the OASIS Executive Director.

841 Copyright © OASIS Open 2005. *All Rights Reserved.*

842 This document and translations of it may be copied and furnished to others, and derivative works
843 that comment on or otherwise explain it or assist in its implementation may be prepared, copied,
844 published and distributed, in whole or in part, without restriction of any kind, provided that the
845 above copyright notice and this paragraph are included on all such copies and derivative works.
846 However, this document itself does not be modified in any way, such as by removing the
847 copyright notice or references to OASIS, except as needed for the purpose of developing OASIS
848 specifications, in which case the procedures for copyrights defined in the OASIS Intellectual
849 Property Rights document must be followed, or as required to translate it into languages other
850 than English.

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

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