



Reference Model for Service Oriented Architectures

Working Draft 10ED, 1 October 2005

Document identifier:

wd-soa-rm-10ED

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 Reference Model for Service Oriented Architectures is an abstract framework for understanding significant entities and relationships between 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 service oriented architectures or in training and explaining SOA. A reference model is not directly tied to any standards, technologies or other concrete implementation details. 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 focuses on the field of software architecture. While the concepts and relationships described may apply to other "service" environments, this specification makes no attempt to completely account for use outside of the software domain.

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.

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:

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

44

45 The errata page for this specification is at:

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

47 **Table of Contents**

48 1 Introduction 4
49 1.1 What is a reference model 4
50 1.2 Audience 4
51 1.3 How to use the reference model 4
52 1.4 Notational Conventions 5
53 1.5 Relationships to Other Standards 5
54 2 Service Oriented Architecture 6
55 2.1 What is SOA? 6
56 3 The Reference Model 8
57 3.1 Overview of model 8
58 3.2 The Reference Model 9
59 3.2.1 Service 9
60 3.2.2 Service description 10
61 3.2.3 Descriptions and Metadata 12
62 3.3 Interacting with services 13
63 3.3.1 Data model 13
64 3.3.2 Behavioral model 14
65 3.3.3 Actualized Services 16
66 3.4 Real World Effect 16
67 Policies and Contracts 17
68 3.5 17
69 3.5.1 Service Policy 17
70 3.6 Service Discoverability 18
71 4 Conformance Guidelines 20
72 5 References 21
73 5.1 Normative 21
74 5.2 Non-Normative 21
75 Appendix A. Glossary 22
76 Appendix B. Acknowledgments 25
77 Appendix C. Notices 26
78

79 1 Introduction

80 The notion of Service Oriented Architecture (SOA) has received significant attention within the
81 software design and development community. The result of this attention is the proliferation of
82 many conflicting definitions of SOA. Whereas SOA architectural patterns (or *reference*
83 *architectures*) may be developed to explain and underpin a generic design template supporting a
84 specific SOA, a reference model is intended to provide an even higher level of commonality, with
85 definitions that should apply to *all* SOA.

86 1.1 What is a reference model

87 A reference model is an abstract framework for understanding significant relationships among the
88 entities of some environment that enables the development of specific architectures using
89 consistent standards or specifications supporting that environment. A reference model consists of
90 a minimal set of unifying concepts, axioms and relationships within a particular problem domain,
91 and is independent of specific standards, technologies, implementations, or other concrete
92 details.

93 The purpose of a reference model is to provide a common conceptual framework that can be
94 used consistently across and between different implementations and is of particular use in
95 modeling specific solutions.

96 The goal of this reference model is to define the essence of service oriented architecture, and
97 emerge with a vocabulary and a common understanding of SOA. It provides a normative
98 reference that remains relevant for SOA as an abstract and powerful model, irrespective of the
99 various and inevitable technology evolutions that will impact SOA.

100 1.2 Audience

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

- 102 • Architects and developers designing, identifying or developing a system based on the
103 service-oriented paradigm.
- 104 • Standards architects and analysts developing specifications that rely on service oriented
105 architecture concepts.
- 106 • Decision makers seeking a "consistent and common" understanding of service oriented
107 architecture.
- 108 • Users who need a better understanding of the concepts and benefits of service oriented
109 architecture.

110 1.3 How to use the reference model

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

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

116 Section two introduces the concept of SOA and identifies some of the ways that it differs from
117 previous paradigms for distributed systems. Section two offers guidance on the basic principles of
118 service oriented architecture. This can be used by non-technical readers to gain an explicit
119 understanding of the core principles of SOA and by architects as guidance for developing specific
120 service oriented architectures.

121

122 Section three introduces the Reference Model for SOA. First, the main axioms, key concepts and
123 relationships between those concepts are introduced followed by more detailed sections on the
124 main concepts, namely *service* is defined along with *service description*. There then follows a
125 section detailing interaction between services, followed by service policies and expectations.
126 Finally, the concept of service discoverability is introduced.

127 Section four addresses compliance with this reference model.

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

130

131 **1.4 Notational Conventions**

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

134

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

136 **1.5 Relationships to Other Standards**

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

- 138 • Considers its work "service oriented";
- 139 • Makes (publicly) an adoption statement to use the Reference Model for SOA of this TC
140 as a base or inspiration for their work; and
- 141 • Standards or technologies that claim to be service oriented.

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

2 Service Oriented Architecture

2.1 What is SOA?

146 Service Oriented Architecture (SOA) is a paradigm for organizing and using distributed
147 capabilities that may be under the control of different ownership domains. It is natural in such a
148 context to think of one person's needs being met by capabilities offered by someone else; or, in
149 the world of distributed computing, one computer agent's requirements being met by a computer
150 agent belonging to a different owner. There is not necessarily a one-to-one correlation between
151 needs and capabilities; the granularity of needs and capabilities vary from fundamental to
152 complex, and any given need may require the combining of numerous capabilities while any
153 single capability may address more than one need. The perceived value of SOA is that it provides
154 a powerful framework for matching needs and capabilities and for combining capabilities to
155 address those needs.

156 Visibility, interaction, and effect are key concepts for describing the SOA paradigm. **Visibility**
157 refers to the capacity for those with needs and those with capabilities to be able to see each
158 other. This is typically done through generating metadata to describe such aspects as functions
159 and technical requirements, related constraints and policies, and mechanisms for access or
160 response. The metadata must be in a form (or can be transformed to a form) in which its syntax
161 and semantics are widely accessible and understandable.

162 Whereas visibility introduces the possibilities for matching needs to capabilities (and vice versa),
163 **interaction** is the activity of sorting through the possibilities to determine which match or
164 combination of matches provides an adequate and feasible response to the needs and then
165 activating that response. Typically mediated by the exchange of messages, the interaction
166 proceeds through a series of information exchanges and invoked actions. There are many facets
167 of interaction; but they are all grounded in a particular **execution context** – the set of technical
168 and business elements that eventually form a path between those with needs and those with
169 capabilities and that permit information to be exchanged, actions to be performed and provides a
170 decision point for any policies and contracts that may be in force.

171 The purpose of using a capability is to realize one or more **real world effects**. At its core, an
172 interaction is “an act” as opposed to “an object” and the result of an interaction is an effect (or a
173 set/series of effects). We are careful to distinguish *public* actions and *private* actions; private
174 actions are inherently unknowable by other parties. On the other hand, public actions result in
175 changes to the *state* that is shared (at least) between those involved in the current execution
176 context. Real world effects are, then, couched in terms of changes to this shared state.

177 The expected effects should be made visible as part of the capability metadata and form an
178 important part of the decision on whether a given capability matches similarly described needs.
179 At the interaction stage, the description of real world effects establishes the expectations of those
180 using the capability. Note, it is not possible to describe every effect from using a capability and,
181 in fact, a cornerstone of SOA is that one using a capability does not need to know all the details.

182 To this point, this description of SOA has yet to mention what is usually considered the central
183 concept: the service. It is likely that both needs and capabilities exist outside of SOA but a major
184 driver for SOA is the perceived improvement in bringing needs and capabilities together. **In SOA,**
185 **services are the mechanism by which needs and capabilities are brought together.** SOA is
186 not the solution of domain problems but rather a way of organizing a wider array of possibilities to
187 generate a domain solution. By itself, SOA does not provide a solution to a difficult domain
188 problem where a satisfactory solution does not already exist. SOA can, however, provide an
189 organizing and delivery paradigm that enables one to get more value from use of both solutions
190 which are locally “owned” and solutions under the control of others. It also enables one to
191 express solutions in a way that makes it easier to modify or evolve the identified solution or to try

192 alternate domain solutions. SOA does not provide any domain elements of a solution that do not
193 exist without SOA.

194 The concepts of visibility, interaction, and effect apply directly to services in the same manner as
195 these were described for the general SOA paradigm. Visibility is established through the **service**
196 **description** which contains the information necessary to interact with the service and describes
197 this in such terms as the service inputs, outputs, and associated semantics. The service
198 description also conveys what is accomplished when the service is invoked and the conditions for
199 invoking the service. In general, entities (people and organizations) offer capabilities through
200 services and act as **service providers**. Those with needs who make use of capabilities through
201 their associated services are referred to as **service consumers**. The service description allows
202 prospective consumers to decide if the service is suitable for their current needs and establish
203 whether a consumer satisfies any requirements of the service provider to be permitted access.

204 Having described what is SOA, it is appropriate to note several things which are related but are
205 not necessary attributes or restrictions.

- 206 • SOA identifies necessary aspects of interactions involving multiple ownership domains;
207 however, it does not directly embody concepts relating to ownership.
- 208 • SOA is commonly implemented using Web services, but services can be made visible,
209 support interaction, and generate effects through other implementations.

210 In most discussions of SOA, the terms “loose coupling” and “coarse-grained” are commonly
211 applied as SOA concepts, but these terms have intentionally not been used in the current
212 discussion because they are subjective and without useful metrics to indicate when the ideals
213 have been accomplished. In terms of needs and capabilities, SOA will be most effective when it
214 is bringing to bear a solution and not the “fine-grained” pieces where those pieces are unlikely to
215 be reusable outside that particular solution. That said, granularity and coarseness are usually
216 relative to detail for the level of the problem being addressed, e.g. one that is more strategic vs.
217 one down to the algorithm level, and defining the optimum level is not amenable to counting the
218 number of interfaces or the number or types of information exchanges connected to an interface.

219 The value of SOA is that it provides a simple scalable paradigm for organizing large networks of
220 systems that require interoperability to realize the value inherent in the individual components.
221 Indeed, SOA is scalable because it makes the fewest possible assumptions, including about the
222 network and also minimizes any trust assumptions that are often implicitly made in smaller scale
223 systems. Loose coupling and coarse-grained as concepts do not address these more
224 fundamental aspects that are critical to SOA success.

225 3 The Reference Model

226 A service oriented architecture represents a uniform means to offer, discover and interact with
227 capabilities to produce desired effects consistent with measurable preconditions and
228 expectations. This section introduces the main concepts and a detailed discussion of the
229 concepts and their relationships are in the sections that follow.

230 3.1 Overview of model

231 A key concept of SOA is that of **service**. In general, entities (people and organizations) create
232 capabilities to solve or support a solution for the problems they face in the course of their
233 business. SOA is a way to organize the world around this key concept of service. The noun
234 "service" is defined in dictionaries as "The performance of work (a function) by one for another."
235 However, service, as a concept, is a unifier for the following related ideas:

236

- 237 • The capability to perform work for another
- 238 • The specification of the work offered for another
- 239 • The offer to perform work for another

240 Therefore, this imprecise term is partitioned into separate, yet interrelated, precise concepts.
241 These concepts are an offer, interaction and effect.

242

243 The concept of an offer follows directly from the dictionary definition of service: 'by one' and 'for
244 another.' In general terms, an offer is a proposal. These offers are made by providers which
245 may possess a capability that address a need. The convergence of a capability and a need
246 results in an interaction. At its core, an interaction is "an act" as opposed to "an object."
247 Therefore, this concept focuses the necessary interfaces and behavior to support the interaction.
248 The final concept focuses on the consequences of interaction.

249 The concepts above emphasize a distinction between a capability and the ability to bring that
250 capability to bear in the context of SOA. It is assumed that the capability was created and exists
251 independently of SOA.

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

257 In order to use a service, it is necessary to know that it exists, what is accomplished if the service
258 is invoked, how the service is invoked, and other characteristics. Collectively this is the service
259 **visibility**. When given an explicit searchable form, this information allows, for example,
260 prospective consumers to decide if the service is suitable for their current needs and establish
261 whether a consumer satisfies any requirements of the service provider to be permitted access.
262 This information constitutes the **service description**.

263 It is often the case that the consumer of a service and the provider of a service belong in different
264 ownership domains. As a result, there is a natural distinction to be drawn between the public
265 interactions with a service and the private actions of both the service provider and consumer. This
266 distinction maintains and encourages independence of each service participant which, in turn,
267 greatly enhances the scalability and security attributes of SOA. Focus can be directed to the
268 public aspects of using a service by examining the **conditions** of using a service and the
269 **expectations** that arise as a result of using the service. Service conditions are loosely associated
270 with the **service policies** and the expectations with **service contracts**.

271

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

276 Finally we identify **discoverability** as a key concept of SOA. Discoverability refers to the
277 possibility and mechanisms by which service consumers and providers can be brought together.
278 There are many possible mechanisms by which discoverability may be achieved. Registries or
279 repositories of service descriptions are undoubtedly powerful means of achieving this, but SOA is
280 not limited to these.

281 SOA identifies necessary aspects of interactions involving multiple ownership domains; however,
282 it does not directly embody concepts relating to ownership.

283 **3.2 The Reference Model**

284 **3.2.1 Service**

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

290

291 A service is invoked through a service interface, where the interface comprises the specifics of
292 how to access the underlying capabilities. There are no constraints on what constitutes the
293 underlying capability or how access is implemented by the service provider. Thus, the service
294 could carry out its described functionality through one or more automated and/or manual
295 processes that themselves could invoke other available services. A service is opaque in that its
296 implementation is typically hidden from the service consumer except for (1) the data model
297 exposed through the published service interface and (2) any information included as metadata to
298 describe aspects of the service which are needed by service consumers to determine whether a
299 given service is appropriate for the consumer's needs. The consequence of invoking a service is
300 an expectation of one or more real world effects. The effects may include:

301

- 302 1. information returned in response to a request,
- 303 2. processing done in response to a request to change the state of defined entities, or
- 304 3. some combination of (1) and (2).

305

306 Note, the user in (1) does not typically know how the information is generated, e.g. whether it is
307 extracted from a database or generated dynamically; in (2), the user does not typically know how
308 the state change is effected. In either case, the service consumer would need to provide input
309 parameters required by the service and the service would return information, status indicators, or
310 error descriptions, where both the input and output are as described by the data model exposed
311 through the published service interface. Note that the service may be invoked without requiring
312 information input from the consumer (other than a command to initiate action) and may
313 accomplish its functions without providing any return or feedback to the consumer.

314 The service concept above has emphasized a distinction between a capability that represents
315 some functionality created to address a problem or a need and the service that forms the point of
316 access to bring that capability to bear in the context of SOA. It is assumed that the capability was
317 created and exists outside of the SOA. One of the major benefits of SOA is enabling the
318 capability to be applicable to an expanded realm of relevant problems. In actual use, maintaining

319 this distinction may not be critical (i.e. the service may be talked about in terms of being the
320 capability) but the separation is pertinent in terms of a clear expression of the nature of SOA and
321 the value it provides.

322 **3.2.2 Service description**

323 The service description represents the information needed in order to use a service. It may be
324 considered part of or the complete set of the metadata (see Section 3.2.3) associated with a
325 service. In any case, the service description overlaps and shares many common properties with
326 service metadata. In most cases, there is no one “right” set of metadata but rather the metadata
327 content depends on the context and the needs of the parties using the associated entity. The
328 same holds for a service description. While there are certain elements that are likely to be part of
329 any service description, most notably the data model, many elements such as function and policy
330 may vary.

331

332 Best practice suggests that the service description should be represented using a standard,
333 reference-able format. Such a format facilitates the use of common processing tools (such as
334 discovery engines) that can capitalize on the service description.

335

336 While the concept of a SOA supports use of a service without the service consumer needing to
337 know the details of the service implementation, the service description makes available critical
338 information that a consumer needs in order to decide whether or not to use a service. In
339 particular, a service consumer must possess the following items of information:

340

- 341 1. That the service exists and is reachable;
- 342 2. That the service performs a certain function or set of functions;
- 343 3. That the service operates under a specified set of constraints and policies;
- 344 4. That the service will (to some implicit or explicit extent) comply with policies as prescribed
345 by the service consumer;
- 346 5. How to interact with the service in order to achieve the required objectives, including the
347 format and content of information exchanged between the service and the consumer and
348 the sequences of information exchange that may be expected.

349

350 Subsequent sections of this document will deal with these aspects of a service in details but the
351 following subsections will describe the relationship of these information items to the service
352 description.

353 **3.2.2.1 Service Reachability**

354 A service description should include sufficient data to permit a service consumer and service
355 provider to exchange information. This might include metadata (such as the location of the
356 service and what information protocols it supports and requires) and whether the service is
357 currently reachable or not.

358 **3.2.2.2 Service Functionality**

359 Item 2 relates to the need to unambiguously express the function(s) of the service and the real
360 world effects (see section 3.5) that result from it being invoked. This portion of the description
361 needs to be expressed in a way that is generally understandable by service consumers but able
362 to accommodate a vocabulary that is sufficiently expressive for the domain for which the service
363 provides its functionality. The description of functionality may include, among other possibilities,
364 a textual description intended for human consumption or identifiers or keywords referenced to

365 specific machine-process-able definitions. For a full description, it may be useful to indicate
366 multiple identifiers or keywords from a number of different collections of definitions.

367

368 Part of the description of functionality may include underlying technical assumptions that
369 determine the limits of functionality exposed by the service or of the underlying capability. For
370 example, the amounts dispensed by an automated teller machine (ATM) are consistent with the
371 assumption that the user is an individual rather than a business. To use the ATM, the user must
372 not only adhere to the policies and satisfy the constraints of the associated financial institution
373 (see Section 3.2.2.3 for how this relates to service description and Section 3.5 for a detailed
374 discussion) but the user is limited to withdrawing certain fixed amounts of cash and a certain
375 number of transactions in a specified period of time. The financial institution, as the underlying
376 capability, does not have these limits but the service interface as exposed to its customers does,
377 consistent with its assumption of the needs of the intended user. If the assumption is not valid,
378 the user may need to use another service to access the capability.

379

380 **3.2.2.3 Policies Related to a Service**

381 Items 3 and 4 from Section 2.2.2 relate to the service description's support for associating
382 constraints and policies with a service and providing necessary information for prospective
383 consumers to evaluate if a service will act in a manner consistent with the consumer's constraints
384 and policies.

385

386 In some situations the consumer may similarly provide an indication of its constraints and policies
387 to support a service's need to do a similar evaluation of suitability. Thus, both prospective
388 consumers and providers are likely to use the service description (and the consumer description)
389 to mutually establish what section 3.3.3 refers to as the *execution context*.
390

391 **3.2.2.4 Service Interface**

392 The service interface is the means referred to in Item 5 for interacting with a service. It includes
393 the specific protocols, commands, and information exchange by which actions are initiated that
394 result in the real world effects as specified through the service functionality portion of the service
395 description.

396

397 The specifics of the interface should be syntactically represented in a standard reference-able
398 format. These prescribe what information needs to be provided to the service in order to exercise
399 its functionality and/or the results of the service invocation to be returned to the service
400 consumer. This logical expression of the set of information items associated with the
401 consumption of the service is often referred to as the service's data model. It should be noted
402 that the particulars of the standard reference-able format is beyond the scope of the reference
403 model. However, requiring that mechanisms be available (in order to define and retrieve such
404 definitions) is fundamental to the SOA concept.

405 While this discussion refers to a standard reference-able syntax for service descriptions, it is not
406 specified how the consumer accesses the interface definition nor how the service itself is
407 accessed. However, it is assumed that for a service to be usable, its interface must be
408 represented in a format that allows interpretation of the interface information by its consumers.

409 **3.2.2.5 An Example of Using Information Contained in the Service** 410 **Description**

411 The following example may help clarify the concepts related to service and service description.
412

413 A utility has the capacity to generate and distribute electricity (the underlying capability). A
414 consumer accesses electricity generated (the service) via a wall outlet (service interface). In order
415 to use the electricity (what is provided by invoking the service), a consumer needs to understand
416 what type of plug to use; the utility presumes that the customer will only connect devices that are
417 compatible with the voltage provided; and the consumer in turn assumes that compatible devices
418 can be connected without damage or harm (service assumptions). A residential or business user
419 will need to open an account with the utility in order to use the supply (service constraint) and the
420 utility will meter usage and expects the consumer to pay for use at the rate prescribed (service
421 policy). When the consumer and utility agree on constraints and polices (service contract), the
422 consumer can receive electricity using the service. Another person (say, a visitor to someone
423 else's house) may use a contracted supply without any relationship with the utility or any
424 requirement to also satisfy the initial service constraint but would nonetheless be expected to be
425 compatible with the service interface. In certain situations (for example, excessive demand), a
426 utility may limit supply or institute rolling blackouts (service policy). A consumer might lodge a
427 formal complaint if this occurred frequently (consumer's implied policy). In this example, the
428 underlying capability would still exist and be usable even if every device were required to be hard-
429 wired to the utility's equipment, but this would result in a very different service and service
430 interface.

431 **3.2.3 Descriptions and Metadata**

432 One of the hallmarks of a Service Oriented Architecture is the degree of documentation and
433 description associated with it; particularly *machine process-able descriptions* – otherwise known
434 as *metadata*.

435 The purpose of this metadata is to facilitate integration, particularly across ownership domains.
436 By providing public descriptions, it makes it possible for potential participants to construct
437 applications that use services and even offer compatible services. Standardizing the formats of
438 such metadata reduces the cost and burden of producing the descriptions necessary to promote
439 reuse and integration.

440 **3.2.3.1 The roles of description**

441 An important additional benefit of metadata – as opposed to informal natural language
442 descriptions – is its potential to facilitate automated software development. Both service providers
443 and service consumers can benefit from such automation – reducing the cost of developing such
444 systems.

445

446 For example, metadata can be used as a basis of discovery in dynamic systems. Metadata can
447 assist in managing a service, validating and auditing usage of services which may also be
448 simplified by rich metadata. It can also help ensure that requirements and expectations
449 (regarding the content of any data interchanged) are properly interpreted and fulfilled.

450 **3.2.3.2 The limits of description**

451 There are well-known theoretic limits on the effectiveness of descriptions – it is simply not
452 possible to specify, completely and unambiguously the precise semantics of a service.

453 There will always be unstated assumptions made by the describer of a service that must be
454 implicitly shared by readers of the description. This applies to machine process-able metadata as
455 well as to human readable documentation.

456 Luckily, complete precision is not necessary either – what is required is sufficient precision to
457 enable required functionality.

458 Another kind of limit of descriptions is more straightforward: describing a service (for example)
459 does not eliminate the requirement for making a choice. For example, a service directory might
460 have the descriptions of many services – provided by many organizations. An automatic search

461 of that directory is therefore likely to return multiple responses to any mechanical search criteria.
462 At some point this set of responses has to be converted into a choice of a single service in order
463 for a service consumer (say) to see the required function performed. In a multi-vendor scenario,
464 that choice must also take into account real world aspects of the service – such as whether the
465 service consumer can identify the provider, can or should trust the provider, and whether the
466 provider is reliable and timely in delivering the service offered. It is unlikely that such factors can
467 be easily and securely encoded in descriptions and search criteria.

468 **3.3 Interacting with services**

469 Interacting with a service involves exchanging information with the service and performing actions
470 against the service. In many cases, this is accomplished by sending and receiving messages, but
471 there are other modes possible that do not involve explicit message transmission. However, for
472 simplicity, we often refer to message exchange as the primary mode of interaction with a service.
473 Together the types of information exchanged and mechanisms used constitute the **service**
474 **interface** – see Section 3.2.2.4.

475 The key concepts that are important in understanding what it is involved in interacting with
476 services are the **data model**, the **process model**, the **execution context** and the **expectations**
477 about the interaction.

478 **3.3.1 Data model**

479 The data model of a service is a characterization of the information that may be exchanged with
480 the service.

481 The scope of the data model includes the format of documents and messages, the structural
482 relationships within those documents and also the definition of terms used within those
483 documents. Only information and data that are potentially exchanged with a service are generally
484 included within that service's data model.

485 Loosely, one might partition the interpretation of an informational block into structure (syntax) and
486 meaning (semantics); although both are part of the data model. Particularly for information that is
487 exchanged across an ownership boundary, an important aspect of the service data model is the
488 interpretation of strings and other tokens in the information.

489 **3.3.1.1 Structure**

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

494 A described data model typically has a great deal to say about the form of messages, about the
495 types of the various components of messages and so on. However, pure “typed” information is
496 not sufficient to completely describe the appropriate interpretation of data.

497 **3.3.1.2 Ontology**

498 The primary task of any communication infrastructure is to facilitate the exchange of information
499 and the exchange of intent. For example, a purchase order combines two somewhat orthogonal
500 aspects: the description of the items being purchased and the fact that one party intends to
501 purchase those items from another party. Even in purely internal exchanges within an Enterprise
502 Application suite, exchanges have similar aspects: this is an update to the customer profile with
503 these changes.

504 With information being exchanged potentially across ownership boundaries a critical issue is the
505 interpretation of the data. This interpretation must be consistent between the participants in the
506 service interaction. Consistency in interpretation of information is a stronger requirement than

507 merely type (or structure) consistency – the tokens in the data itself must also have a shared
508 basis.

509 For example, within a street address structure, there is a huge potential for variability in
510 representing addresses. For example, an address in San Francisco, California may have
511 variations in the way the city is represented: SF, San Francisco, San Fran, the City by the Bay
512 are all alternate denotations of the same city. For successful exchange of address information, all
513 the participants must have a consistent view of the meaning of the address tokens. How a city
514 name is interpreted is a simple but classic problem in ensuring that addresses can be reliably
515 shared.

516 An ontology is a formal description of terms and the relationships between them in a given
517 context. Most commonly, the relationships are class relationships – one term represents a
518 concept that is a sub-class of another. However, relationships are not limited to the sub-class
519 relationships; other aspects of concepts can also be usefully represented; such as the range of
520 possible values given property can take and whether the property is functional or not.

521 The role of explicit ontologies is to provide a firm basis for selecting correct interpretations for
522 tokens in messages. For example, an ontology can be used to capture the alternate ways of
523 expressing the names of cities as well as distinguishing a city name from a street name.

524 Ontologies also provide a point of context to facilitate the *reinterpretation* of data – for example
525 that a 3/8" steel washer may be a potential replacement for a 1cm spacer. Such a reinterpretation
526 is effectively represented as a particular traversal of the graph of concepts and relationships
527 embodied in the ontology. How much automation of ontology walking is appropriate will depend
528 on the nature of the service and the service participants.

529 Note that, for the most part, it is not expected that service consumers and providers would
530 actually exchange ontologies in their interaction – the role of the ontology is a background one – it
531 facilitates sound interactions. Hence ontology references are mostly to be found in **service**
532 **descriptions**.

533 More specifically, and in order for a service to be consistent, the service should make consistent
534 use of terms as defined in an ontology. Specific domain semantics are beyond the scope of this
535 reference model; but there is a requirement that the service interface enable providers and
536 consumers to identify unambiguously those definitions that are relevant to their respective
537 domains.

538 **3.3.2 Behavioral model**

539 The second key requirement for successful interactions with services is knowledge of the actions
540 invoked against the service and the process or temporal aspects of interacting with the service.
541 Loosely, this can be characterized as knowledge of the actions on, responses to and temporal
542 dependencies between actions on the service.

543 For example, in a security-controlled access to a database, the actions available to a service
544 consumer include presenting credentials, requesting database updates and reading results of
545 queries. The security may be based on a challenge-response protocol. For example, the initiator
546 presents an initial token of identity, the responder presents a challenge and the initiator responds
547 to the challenge in a way that satisfies the database. Only after the user's credentials have been
548 verified will the actions that relate to database update and query be accepted.

549 The sequences of actions involved are a critical aspect of the knowledge required for successful
550 use of the secured database.

551 There are other aspects of the behavior of services that are important. These include, for
552 example, whether the service is transactional, idempotent or long running. As a particular
553 example, a service that supports updating an account balance with a transaction is typically
554 idempotent; i.e., the state of the account would not be affected should a subsequent interaction
555 be attempted for the same transaction.

556 3.3.2.1 Action model

557 The **action** model of a service is about the individual actions that may be invoked against the
558 service. Of course, a great portion of the behavior resulting from an action may be private;
559 however, the expected public view of a service surely includes the implied effects of actions.
560

561 For example, in a service that manages a bank account's state, it is not sufficient to know that to
562 use the service you need to exchange a given message (with appropriate authentication tokens).
563 It is also of the essence that using the service may actually affect the bank account – withdrawing
564 cash from it for example.

565 3.3.2.2 Process Model

566 The **process model** characterizes the temporal relationships between actions and other higher-
567 order attributes of interacting with the service. It is fairly common to partition the process model
568 associated with a service into two levels: the particular sequences of operations needed to
569 achieve single service exchanges and longer term transactions. These two levels may be nested
570 – a long running transaction is often composed of sequences of exchange patterns.

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

576 3.3.2.3 Higher-order attributes of processes

577 Beyond the straightforward mechanics of interacting with a service there are other, higher-order,
578 attributes of services' process models that are also often important. These can include whether
579 the service is **idempotent**, whether the service is **long-running** in nature and whether it is
580 important to account for any **transactional** aspects of the service.

581 A service is idempotent if subsequent attempts to perform identical transactions are discounted.
582 For example, it is often important that a bank will only cash a cheque once – subsequent attempts
583 to cash the same cheque should be ignored or rejected. Note that idempotency is not the same
584 as effect-free or stateless: a service that always returns the same results is idempotent, but only
585 by virtue of the fact that it does not change from invocation to invocation.

586 Idempotency is an important attribute of a service in an environment where there is a significant
587 possibility that the interaction between the service provider and consumer may be interrupted –
588 whether by a network issue or simply one of the parties dropping out.

589 A service is long-running if the activities engendered by an interaction are likely to persist beyond
590 the immediate interaction itself. For example, a classic book selling service might be viewed as a
591 long-running service: the activity started by the purchase of the book may take days or weeks to
592 complete. It can be important to account for a long-running process as it has implications for the
593 kinds of infrastructure needed – both by the service provider and by the service consumer – in
594 order to be able to keep track of the progress of the interaction.

595 Often, once a business-level contract has been agreed on, it can be difficult or impossible to
596 simply cancel the consequences of the agreement. This is particularly an issue when the
597 agreement of several parties is necessary simultaneously. For example, booking a vacation may
598 require a flight ticket as well as a hotel room – without either the vacation is a bust. However, the
599 airline typically will not have a relationship with the hotel; so if it turns out that there are no hotel
600 rooms available for the vacation the airline ticket will need to be canceled.

601 The process of reversing a previously completed transaction – backing out of the airline booking
602 for example – is likely to be quite different to the process for the original transaction. Knowledge
603 such compensatory actions is a key aspect of interacting with transactional services.

604 3.3.3 Actualized Services

605 The **execution context** of a service interaction is the set of infrastructure elements, process
606 entities and policy assertions that are deployed as part of the instantiated service interaction. In
607 effect, the execution context defines the point of contact between abstractions such as service
608 descriptions which are mostly about the potential for interaction and an actually executing service.

609 The most basic aspect of an execution context must be consistency between what is actually
610 happening and any descriptions of what ought to be happening. The execution context is the
611 point of measurement between the service description and reality.

612 The execution context is not limited to one side of the interaction; rather it refers to the totality of
613 the interaction – including the service provider, the service consumer and the common
614 infrastructure needed to mediate the interaction. Thus it represents a path between the intentions
615 of the service consumer and those of the service provider.

616 The execution context is a critical touchstone for many aspects of a service interaction – for
617 example, it defines the decision point for any policy enforcement relating to the service
618 interaction. Note that a policy decision point is not necessarily the same as an enforcement point:
619 an execution context is not by itself something that lends itself to enforcement. On the other hand,
620 any enforcement mechanism of a policy is likely to take into account the particulars of the actual
621 service interaction.

622 The execution context also allows us to distinguish services from one another. Different instances
623 of the same service – denoting interactions between a given service provider and different service
624 consumers for example – are distinguished by virtue of the fact that their execution contexts are
625 different.

626 Finally, the execution context is also the context in which the interpretation of data that is
627 exchanged takes place – it is where the *symbol grounding* happens as it were. A particular string
628 has a particular meaning in a service interaction in a particular context – the execution context.

629 3.4 Real World Effect

630 There is always a particular purpose associated with interacting with a service; conversely, a
631 service provider (and consumer) often has a priori conditions that apply to its interactions. The
632 service consumer is trying to achieve some result by interacting with the service, as is the service
633 provider. At first sight, such a goal can often be expressed as “trying to get the service to do
634 something”. This is sometimes known as the **real world effect** of using a service. For example,
635 an airline reservation service can be used in order to book travel – the desired real world effect
636 being, presumably, a seat on the right airplane.

637 The internal actions that a service providers and consumers perform as a result of participation in
638 service interactions are, by definition, private and fundamentally unknowable.¹ By unknowable we
639 mean both that external parties cannot see others’ private actions and, furthermore, should not
640 have explicit knowledge of them. Instead we focus on the state that is shared between the parties

¹ A similar analysis applies to service consumers: just how a consumer of a service decides which requests to make and which actions to performs is something that the service provider cannot determine.

641 – the **shared state**. Actions by service providers and consumers lead to modifications of this
642 shared state; and that in turn leads to modified **expectations** by the participants.

643 When an airline has confirmed a seat for a passenger on a flight this represents a fact that both
644 the airline and the passenger share – it is part of their shared state. Thus the real world effect of
645 booking the flight is a modification of this shared state – the existence of the booking. Flowing
646 from the shared facts, both the passenger, the airline and interested third parties may make
647 inferences – in particular that when the passenger arrives at the airport the airline confirms the
648 booking and permits the passenger onto the airplane (subject of course to the passenger meeting
649 the other requirements for traveling).

650 Of course, in order for the airline to know that the seat is confirmed it will likely invoke some kind
651 of private action to record the reservation; but, by minimizing assumptions about how the airline
652 fulfils its contracts, potential for smooth interoperation is maximized. Such minimization principles
653 represent a key success factor for scalability.

654 Although there is not necessarily a one-to-one correspondence, the natural container for the
655 conditions applying to a service is the **service policy**. Similarly, the natural container for the
656 expectations arising from a service is the **service contract**.

657 **3.5 Policies and Contracts**

658 Broadly speaking, a **policy** represents some form of constraint or condition on the use,
659 deployment or description of an owned entity. A **contract**, however, represents an agreement of
660 some kind; often the agreement is also about the conditions of use of a service and what the
661 expected real world effects might be. The reference model is focused primarily on the concept of
662 policies and contracts as they apply to services.

663 **3.5.1 Service Policy**

664 In abstract, a policy is a statement of the obligations, constraints or other conditions of use of a
665 given service that expresses intent on the part of a participant. More particularly, policies are a
666 way for expressing the relationship between the **execution context** and the **data** and **behavior**
667 **models** associated with the service.

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

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

675 A policy always represents a participant’s point of view. An assertion becomes the policy of a
676 participant when they make it their policy – this linking is normally not part of the assertion itself.
677 For example, if the service consumer declares that “All messages are triple-DES encrypted”, then
678 that reflects the policy of the service consumer. This policy is one that may be asserted by the
679 service consumer independently of any agreement from the service provider.

680 Finally, a policy may be enforced. Techniques for the enforcement of policies depend on the
681 nature of the policy. From a conceptual point of view, service policy enforcement amounts to
682 ensuring that the policy assertion is consistent with the real world. This might mean preventing
683 unauthorized actions to be performed or states to be entered into; it can also mean initiating
684 compensatory actions when a policy violation has been detected. An unenforceable policy is not
685 a policy; it would be better described as a wish.

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

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

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

696 A natural point of contact between service participants and policies associated with the service is
697 in the service description – see Section 3.2.2. It would be natural for the service description to
698 contain references to the policies associated with the service.

699 **3.5.1.1 Service Contract**

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

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

708 Thus, following the analysis above, a service contract is a measurable assertion that governs the
709 requirements and expectations of two or more parties. Unlike policy enforcement, which is
710 usually the responsibility of the policy owner, contract enforcement may involve resolving
711 disputes between the parties to the contract. The resolution of such disputes may involve appeals
712 to higher authorities.

713 Like policies, contracts may be expressed in a machine processable form. Where a contract is
714 used to codify the results of a service interaction, it is good practice to represent it in a machine
715 processable form. This facilitates automatic service composition, for example. Where a contract is
716 used to describe over-arching agreements between service providers and consumers, then the
717 priority is likely to make such contracts readable by people.

718 **3.6 Service Discoverability**

719 A key concept of the SOA Reference Model is the discoverability of services. Discoverability is an
720 important aspect of bringing together the service provider and consumer: A service provider must
721 be capable of making details of the service (notably service description and policies) available to
722 potential consumers; and customers must be capable of finding that information.

723 This may (and commonly does) involve a service provider entering the service description into a
724 service registry and the service consumer searching for an appropriate match to their needs. The
725 SOA concept of discoverability is not restricted to any single mechanism. In some architectures
726 there may not be a registry .

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

732 For example, a service provider may advertise and promote their service by either including it in a
733 service directory or broadcasting it to all consumers; potential consumers may broadcast their
734 particular service needs in the hope that a suitable service responds with a proposal or offer or a
735 service consumer might also “probe()” an entire network to determine if suitable services exist.
736 When the demand for a service is higher than the supply, then by advertising their needs,

737 potential consumers are likely to be more effective than service providers advertising offered
738 services.

739 One way or another, the potential consumer must acquire a sufficient description to evaluate
740 whether the service matches their expectations and, if so, the method for the consumer to
741 establish a contract and invoke the service.

742 Specific SOA reference architectures and implementations will prescribe the mechanisms for
743 actual service discovery, ensuring a service's presence and availability, and failure conditions and
744 error handling.

745

4 Conformance Guidelines

746

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

747

748

749

750

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

751

752

753

- Have entities that can be identified as services as defined by this Reference Model,

754

- Such entities will have descriptions associated with them,

755

- Service entities will have identifiable interaction models, including models of the information exchanged by the services and the temporal behavior of the services

756

757

- It should be possible to identify a means by which consumers of services and providers of services are able to engage; and

758

759

- That there will be identifiable aspects of service entities that correspond to the policies relating to the conditions of use of services and to the expectations that result from interacting with services.

760

761

762

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

763

764

765

766 **5 References**

767 **5.1 Normative**

768 [RFC2119] S. Bradner, *Key words for use in RFCs to Indicate Requirement Levels*,
769 <http://www.ietf.org/rfc/rfc2119.txt>, IETF RFC 2119, March 1997.
770

771 **5.2 Non-Normative**

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

774 **Appendix A. Glossary**

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

778

779 Advertising (or Announcement of Availability)

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

782

783 Agent (requester or provider)

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

785

786 Architecture

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

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

792

793 Authentication

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

796

797 Capability

798 Response to a need that is characterized by a set of preconditions and where the solution is
799 consistent with the expectations of real world effects that would result from making use of the
800 capability.

801

802 (Service) Consumer

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

804

805 Contract

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

807

808 Data Model

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

811

812 Discovery

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

814

815 Encapsulation
816 The act of hiding internal specifications of an entity from the user of that entity, in such a way that
817 the internal data and methods of the entity can be changed without changing the manner in which
818 the entity is used. What is seen by the user is only an interface, or service.
819
820 Framework
821 A set of assumptions, concepts, values, and practices that constitutes a way of viewing the
822 current environment.
823
824 Interface
825 A named set of operations that characterize the behavior of an entity.
826
827 Mediation
828 The transformation, routing, validation and processing of messages.
829
830 Message
831 A serialized set of data that is used to convey a request or response from one party to another.
832
833 Metadata
834 A set of properties of a given entity which are intended to describe and/or indicate the nature and
835 purpose of the entity and/or its relationship with others.
836
837 Negotiation
838 A process that seeks to establish an acceptable basis for a contract between agents for the
839 provision of a service.
840
841 Ontology
842 Represents an agreement within a specific environment of the meanings to be associated with
843 different concepts and their relations to each other.
844
845 Opaqueness
846 The extent to which an agent is able to interact successfully with a service without detecting how
847 the service is implemented.
848
849 Policy
850 A statement of obligations, constraints or other conditions of use of a given service. When a
851 specific set of entities accept such a policy, a contract is usually established.
852
853 Reference Model
854 A reference model is an abstract framework for understanding significant relationships among the
855 entities of some environment that enables the development of specific architectures using
856 consistent standards or specifications supporting that environment.
857 A reference model is based on a small number of unifying concepts. A reference model is not
858 directly tied to any standards, technologies or other concrete implementation details, but it does

859 seek to provide a common semantics that can be used unambiguously across and between
860 different implementations.

861

862 (Service) Requester or provider

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

864

865 Security

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

870

871 Semantics

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

874

875 Service

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

878

879 Service description

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

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

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

884 - the procedures and constraints governing invocation of the service,
885 and thus determine whether the service meets the expectations and requirements of the
886 consumer. Acceptance of the service description by a consumer does not of itself imply a contract
887 to use the service.

888

889 Service Oriented Architecture (SOA)

890 A software architecture of services, policies, practices and frameworks in which components can
891 be reused and repurposed rapidly in order to achieve shared and new functionality. This enables
892 rapid and economical implementation in response to new requirements thus ensuring that
893 services respond to perceived user needs.

894 SOA uses the object-oriented principle of encapsulation in which entities are accessible only
895 through interfaces and where those entities are connected by well-defined interface agreements
896 or contracts.

897

898

899

900

901 **Appendix B. Acknowledgments**

902 The following individuals were members of the committee during the development of this
903 specification:

904 [TODO: insert cte. Members]

905

Appendix C. Notices

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

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

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

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

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

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