



Service Oriented Architecture Reference Model

Working Draft 07, 12 May 2005

Document identifier:

wd-soa-rm-07

Location:

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

Editors:

C. Matthew MacKenzie, Adobe Systems Incorporated, mattm@adobe.com
John Harby, Individual, jharby@gmail.com
Michael Ruiz, BAE Systems PLC, Michael.ruiz@baesystems.com
Christopher Bashioum, Mitre Corporation, cbashioum@mitre.org
Ken Laskey, Mitre Corporation, klaskey@mitre.org
Wesley McGregor, Government of Canada (PWGSC), McGregor.Wesley@tbs-sct.gc.ca
Francis McCabe, Fujitsu Limited, fgm@fla.fujitsu.com
Don Flinn, Individual, flinn@alum.mit.edu
Peter Brown, Individual, peter@justbrown.net
Vikas Deolaliker, Sonoa Systems, Inc., vikas@sonoasystems.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 system a broad variety of applications, this reference model scopes itself to the field of software architecture.

Status:

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

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

42
43 For information on whether any patents have been disclosed that may be essential to
44 implementing this specification, and any offers of patent licensing terms, please refer to
45 the Intellectual Property Rights section of the SOA-RM TC web page at:

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

47
48 The errata page for this specification is at:

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

DRAFT

50 Table of Contents

51	1	Introduction	5
52	1.1	Audience	5
53	1.2	How to Use the Reference Model	5
54	1.3	Notational Conventions	6
55	1.4	Relationships to Other Standards	6
56	2	The Reference Model	8
57	2.1	Services.....	9
58	2.1.1	Service Composition	9
59	2.1.2	Service interface	10
60	2.1.3	Service description.....	11
61	2.2	Policies and contracts	15
62	2.2.1	Service Policy	15
63	2.2.2	Service Contract	16
64	2.3	Semantics.....	17
65	2.3.1	The layers of a semantics of service	17
66	2.3.2	Metadata	18
67	2.3.3	Vocabulary	18
68	2.3.4	Context.....	18
69	2.3.5	Autonomy.....	19
70	2.3.6	Data/Information Model	19
71	2.4	Discovery, Presence and Availability	19
72	2.4.1	Discovery	19
73	2.4.2	Structured vs. Unstructured Discovery	20
74	2.4.3	Service Fabric Constraints.....	20
75	2.4.4	Discovery Methods	20
76	2.4.5	Identification, Understanding and Matching	22
77	2.4.6	Presence and Availability.....	22
78	3	Conformance Guidelines	24
79	4	References.....	25
80	4.1	Normative	25
81	4.2	Non-Normative	25
82		Appendix A. Glossary	26
83		Appendix B. Use Cases and Examples (Non-Normative)	30
84	1	Introduction	31
85	2	Use-Cases	32
86	2.1	Simple SOA.....	32

87	2.2 Intermediate SOA.....	33
88	2.3 Complex SOA.....	35
89	Appendix C. Metadata in the context of a SOA.....	39
90	Appendix D. Acknowledgments.....	41
91	Appendix E. Notices	42
92		

DRAFT

93 1 Introduction

94 The service-oriented architecture (SOA) paradigm has received significant attention within the
95 software design and development industry in recent times resulting in many conflicting definitions
96 of service-oriented architecture. The goal of this reference model document is to define the
97 essence of the service oriented architecture paradigm, and emerge with a vocabulary and a
98 common understanding of SOA.

99

100 This document explicitly avoids defining implementation detail, as doing so would unnecessarily
101 constrain and date the reference model. The goal is to provide a document that can stay relevant
102 through the various technology evolutions that we experience in this industry.

103

104 A reference model cannot be implemented, nor should it be. A reference model is a foundational
105 work that can and should be used to develop architectural patterns and promote effective
106 discourse on derived works.

107 1.1 Audience

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

109

- 110 • Architects and developers designing, identifying or developing a system based on the
111 service-oriented paradigm.
- 112 • Standards architects / analysts developing specifications that relates to or makes use of
113 the service-oriented paradigm.
- 114 • Chief Information Officers and other decision makers seeking a "consistent and common"
115 understanding of service oriented architecture.

116

117 1.2 How to Use the Reference Model

118

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

121

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

125

126 Section 2 introduces the service oriented reference model. First, services are defined and service
127 composition and description are described. A brief overview of the policy components and their

128 relationships is given. This section is provided for the benefit of multiple audiences. Non-technical
129 readers may use this section to gain an explicit understanding of the core principles of SOA.

130

131 Architects are encouraged to use this section as guidance for developing specific service oriented
132 architectures. Section 2 and its subsections are designed to provide guidance for consistent
133 logical divisions of components within architectures. It also helps architects adhere to the basic
134 principles of service-oriented design.

135

136 Section 3 aims to provide guidelines for conformance with the reference model and is aimed at
137 those who wish to explicitly state that their architectures are conformant with this reference
138 model.

139

140 Section 4 provides references to external material used in the reference model.

141

142 The appendices provide several non-normative examples and a glossary to provide clarity of
143 terms whose use may otherwise be ambiguous.

144

145 **1.3 Notational Conventions**

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

148

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

150 **1.4 Relationships to Other Standards**

151

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

153

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

157

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

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

162

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

DRAFT

2 The Reference Model

166

167 Figure 2-1 - SOA Architectural Model introduces the core Service Oriented Architecture (SOA)
168 reference model and its high level components.

169 A service, the fundamental element of a SOA, is decomposed into four distinct aspects and two
170 cross cutting concerns. The four distinct aspects include:

- 171 • Descriptor
- 172 • Policy
- 173 • Contract
- 174 • Data Model

175

176 The “cross cutting” concerns are:

177

- 178 • Semantics
- 179 • Discovery, Presence and Availability

180

181 The Descriptor is comprised of metadata that articulates the interface of a service in order for
182 Service Consumer to understand the service’s externally accessible functionality.

183

184 A Policy is a set of assertions that must be adhered to when a service is invoked.

185


186 A Contract is implied when a Service Consumer makes and invocation request to a service, in
187 substantial alignment with the Policy declaration.

188

189 Data Model is the abstract paradigm used in the invocation and consumption of a Service. A
190 Data Model will likely manifest itself within a concrete architecture as a set of concrete Messages.

191 The cross cutting concerns are defined as aspects that cross several other elements within the
192 object model. Within the SOA-RM, these crosscutting concerns are a Semantic aspect and a
193 Discovery aspect. Semantic agreement on what entities mean with respect to their roles in a
194 system is necessary for service oriented architecture. Many of the components (Service
195 Descriptions, Policies, Contracts and Data Models) need to be available for discovery by potential
196 service consumers to determine both the suitability of a service and their ability to invoke and/or
197 consume the service. The concept of Discovery is to gain awareness of the Presence of the
198 elements and details of their availability.

199



200

201 *Figure 2-1 - SOA Architectural Model*

202

203 **2.1 Services**

204

205 A service is a set of functionality provided by one entity for the use of others. It is invoked
 206 through a software interface but with no constraints on how the functionality is implemented by
 207 the providing entity. Thus, the service could carry out its described functionality through one or
 208 more automated and/or manual processes that themselves could invoke other available services.
 209 A service is opaque in that its implementation is hidden from the service consumer except for (1)
 210 the data model exposed through the published service interface and (2) any information included
 211 as metadata to describe aspects of the service which are needed by service consumers to
 212 determine whether a given service is appropriate for the consumer's needs. Thus while service
 213 opacity is an essential of SOA, it is not absolute.

214

215 **2.1.1 Service Composition**

216

217 Consistent with the axiom of opacity, a Service Consumer cannot see anything behind the service
 218 interface and does not know if one service is actually consuming and aggregating additional other
 219 services. Whether a Service's functions are mapped to a set of classes in some native language
 220 or another service is not important or relevant as far as invoking the service is concerned.

221



12 May 2005

222

223

224 Figure 2 - Service Composition

225

226 Examining Figure 2 - Service Composition above, the service function (for service A) is described
227 in the service description specific to that service. If completing the function depends on two or
228 more serial or parallel paths of execution successfully completing behind the service interface
229 (like calling services B and C) within a certain time frame, that is typically not relevant to state in
230 the service description for service A. Ideally, the service consumer is only concerned with the
231 service's ultimate success or failure. Mapping the functionality to success and failure is the
232 responsibility of the service provider.

233

234 As part of hiding its implementation, when a service is invoked by multiple users in a manner
235 such that a new service invocation is requested before a previous service invocation is
236 completed, the mechanisms the service uses to handle the overlapping (and possibly
237 simultaneous) invocations is not typically revealed to the requester. Indeed, the service may
238 make use of other services providing specialized functionality to support such needs. However,
239 there may be situations, such as quality of service requirements, where the effects of
240 implementation choices have consequences that impact descriptive quantities included in the
241 service metadata. Here, while the implementation details may not be specifically revealed,
242 information derivable from these details will be available to the requester.

243

244 **2.1.2 Service interface**

245

246 The service interface specifies how to access the service and syntactically represents this
247 information in a standard, reference-able format. It prescribes what information needs to be
248 provided to the service in order to exercise its functionality and/or the results of the service
249 invocation to be returned to the service requester. This logical expression of the set of
250 information items associated with the consumption of the service is often referred to as the
251 service's data model. Note, that the service may be invoked without requiring input from the
252 requester and may accomplish its functions without providing any return or feedback to the
253 requester.

254

255 In addition to conforming to a standard, reference-able syntax, the service interface must also
256 make consistent use of SOA semantics as defined in this reference model. This may be
257 represented as a mapping between SOA semantics and the chosen interface syntax. Note, the
258 specific domain semantics of the service provider and service consumer are beyond the scope of
259 this reference model but the reference model does [DOES IT?] address the need for the service
260 interface to enable providers and consumers to unambiguously identify relevant definitions for
261 their respective domains. See detailed discussion of SOA semantics in section 2.3.

262

263 2.1.3 Service description

264

265 As discussed above, the concept of a SOA is based on the use of a service without the service
266 consumer needing to know the details of the service implementation. Hidden details could
267 include the specific logic applied, the mechanism for encoding the logic, or the physical means by
268 which the service is hosted. However to use a service, a service consumer must know

269

- 270 1. The service exists and is available
- 271 2. The service performs a certain function or set of functions
- 272 3. The service operates under a specified set of assumptions, constraints, and policies
- 273 4. The service can be invoked through a specified means, including inputs that the service
274 requires and outputs that will form the response to the invocation.

275

276 The mechanisms to establish presence and availability are discussed elsewhere in this
277 document; items (2) through (4) form the service description. The service interface, as described
278 above, describes the basics of the required inputs and outputs that make up the data model for
279 item (4). The description of functionality and assumptions, constraints, and policies are less
280 specific and more dependent on the context to which the service provider and consumer are
281 aligned.

282

283 The most difficult of the description items is item (2), that capturing service functionality. This
284 aspect of description needs to be expressed in a way that is generally understandable by service
285 consumers but able to accommodate a vocabulary that is sufficiently expressive for the domain
286 for which the service provides its functionality. This may include, among other possibilities, a
287 textual description intended for human consumption or identifiers or keywords referenced to
288 specific machine-process-able definitions. The specification of a single description vocabulary is
289 not only beyond the scope of this reference model; it is unlikely that such a single, general-
290 purpose vocabulary exists or can be developed.

291

292 Assumptions, constraints, and policies are particular descriptive aspects of a service that control
293 if and under what circumstances the service is appropriate or accessible for use. While the line
294 between each of these is vague, the common requirement is that they must be expressed in such
295 a way as to enable corresponding instances to be processed in a consistent, logical fashion. In
296 essence, assumptions, constraints, and policies not only provide information but also the
297 elements of a logical framework that can be interpreted and enforced.

298

299 Assumptions in the technical sense provide conditions that underlie the derivation of the service
300 functionality. For example, a service that calculates the pressure distribution around a body
301 might indicate whether the solution assumes compressible or incompressible flow and whether
302 shock conditions fall within the service capabilities. The appropriate service would be different for
303 a submarine vs. a small passenger aircraft vs. a jet flying at supersonic speeds. This example
304 not only highlights the need to evaluate a set of assumptions but also the need to express
305 assumptions specific to a particular domain of discourse.

306

307 Constraints, like assumptions, can restrict how a service is to be used. While it may not be an
308 underlying assumption, it could be a precondition to a service being accessed. For example, a
309 constraint could be that a prospective consumer needs to prove there is a paid subscription
310 before the service can be accessed.

311

312 Policies express a set of assertions and obligations to which service providers and/or consumers
313 must adhere. To make use of policies, these must be expressed in a way that characteristics of
314 the provider or consumer can be identified to evaluate whether the policy conditions are being
315 satisfied. For example, if policy states that employee salary information can only be accessed by
316 their direct supervisor or the group's designated HR representative, then the required conditions
317 must be visible to and identifiable by the policy evaluation mechanism. There may also be a
318 need to capture the results of policy evaluations and such results may be appropriately included
319 as part of the metadata of the service or the participating entities. Metadata is discussed below
320 and policy is more fully discussed in Policies and contracts.

321

322 **2.1.3.1 The relationship between service description and service metadata**

323

324 The service description may be considered part of or the complete set of the metadata associated
325 with a service (see Appendix META for a discussion of metadata in the context of a SOA) but in
326 any case, the service description overlaps and shares many common properties with service
327 metadata. As noted in Appendix META, there is no one "right" set of metadata but rather the
328 metadata content depends on the context and the needs of the parties using the associated
329 entity. The same holds for a service description. While there are certain elements that are likely
330 to be part of any service description, most notably the data model, many elements such as
331 assumptions and policy may vary. However, the mechanisms to specify the service description
332 should follow a standard, reference-able format that can accommodate the necessary variations
333 and lend themselves to common processing tools (such as discovery engines) to manipulate the
334 service description.

335

336 Consider, for example, the descriptive elements that may apply for a data resource vs. a
337 processing resource. Here, we will assume the resources to be distinguished as follows:

338 A data resource is a source of content that accepts a request and returns a value or set of values
339 in response. The return can be an entity (such as a particular schema), an attribute of an entity
340 (such as when the schema was last modified), or any numerical or textual value or set of values.
341 The content can be static objects stored in some repository or dynamically generated through the
342 use of a processing resource.

343 A processing resource is one that accepts a task and return a status indicating the extent to
344 which the task was completed and information on how the state of entities changed as a result of
345 the processing. One or more processing resources may be invoked as part of a process of
346 submitting a query and being returned a response. From the standpoint of a user (either human
347 or machine), it is unimportant what combination of data and processing resources are invoked as
348 long as the request is satisfied.

349

350 Both types of resources are likely to have items such as a name, a textual description, and
351 possibly a set of descriptors/keywords with a pointer to the vocabulary definition from which the
352 descriptors/keywords are taken. Both resources may also identify responsible parties, including
353 who is responsible for operations, who is responsible for design, who is responsible for
354 implementation

355

356 The description of services to establish their discovery, presence, and availability is imperative for
357 a service-oriented architecture. From a metadata standpoint, there is no significant difference in
358 describing a service that may be considered integral to the SOA infrastructure as opposed to one
359 defined by a specialized participating community. The metadata must support (1) discovery by a
360 user looking for a service to compose a solution, (2) mutual evaluation by the service and the
361 prospective user to decide if service authorization requirements are met by the user and
362 usability/applicability requirements are met by the service, and (3) access/invocation after service
363 and user have mutually satisfy their conditions for use. A common metadata set includes familiar
364 elements such as name, description, and keywords. Access/invocation and pedigree are
365 included per their descriptions above, and security and SLA metadata, while not fully analyzed
366 but tentatively identified as Upper Level metadata, are included in the notional elements.

367

368 The Service metadata does include several unique elements. Two instances of Responsible
369 Party metadata are used: one to identify the entity that is responsible for the design,
370 development, and maintenance of the software that comprises the service; a second entity is
371 identified who is responsible for service operation issues for NCES users. Both instances of
372 Responsible Party will likely use the Person/Organization or Title/Position building blocks.

373

374 Service metadata also contains Version and Status elements. Both of these should reference
375 documentation that defines the values which are applied to these elements. The version
376 numbering may follow a format specified through NCES governance, but a general requirement
377 to include a pointer to a defining document supports the current directed use, modifications to the
378 directed use, and any other versioning algorithm that the development community finds useful.
379 The Status element is intended to reflect the status as determined by the developer (e.g., current
380 version, beta, superseded), and may be seen as a counterpart to pedigree which is an evaluation
381 from the users. As with version numbering, NCES governance could specify authoritative status
382 states, but the reference to a defining document support both effective governance and future
383 contingencies.

384

385 It is assumed that there may be multiple instances of the Access/Invocation metadata bin to
386 expose different aspects of a service. For example, the access may be different depending on
387 the guaranteed quality of service. Recall that the Access/Invocation metadata bin included a list
388 of pre-qualified users as a notional means of speeding access without repeating all the
389 authorization checks previously satisfied for that access point. For a service with multiple access
390 points, the Service metadata includes the notional capability to provide a global prequalification
391 list that alleviates the need to provide a duplicate list with each Access/Invocation metadata
392 instance. It is expected that any differences in the local list would override prequalification in the
393 global list.

394

395 **2.1.3.2 Application metadata**

396 In the AoA analysis, application are considered processing resources that a service is designed
397 to invoke. For example, a Discovery service could access several discovery engines and
398 Application metadata would provide the basis on which one would choose a particular application
399 and how that applications would be executed. Structural metadata detailing an application API
400 might be used by a service developer in developing service access to the application.

401

402 The purpose of Application metadata is to support cataloguing of applications that were not
403 originally intended for Web service access. In most ways, the information included in the
404 Application metadata is the same as identified for a service, and most of the Service metadata
405 discussion is applicable here. The important point is again that many applications may provide
406 similar capabilities and may even be alternatives to be invoked by a single service. It is the
407 purpose of the metadata to allow the user to discriminate among the alternatives and invoke the
408 one that is best suited for the current tasking needs.

409

410 **2.1.3.3 Data Source metadata**

411

412 Data Source metadata provides the counterpoint to Service metadata. The metadata must
413 support (1) discovery by a user needing data, (2) mutual evaluation by the data source and the
414 prospective user to decide if data source authorization requirements are met by the user and user
415 usability/applicability requirements are met by the data source, and (3) read and write, as
416 appropriate, after the data source and user have mutually satisfy their conditions for use. The
417 data source does not have to be a database; in fact, the user may have no specific knowledge of
418 how the data is stored or accessed. Any information needed to evaluate the applicability of the
419 data source should be supported via constraint descriptions and established through pedigree
420 evaluation.

421

422 Most of the notional elements for Data Source metadata are the same as described for Service
423 metadata, and those explanations also apply here. The one notional element added was update
424 cycle. This may not be applicable to all data sources. For example, a database tracking truck
425 parts is likely to be continually updated and it is impractical and of questionable use to update the
426 data source metadata every few minutes. However, simply indicating the refresh cycle is
427 continuous may be of use. For a data mart, the refresh cycle is more relevant because it
428 indicates whether the contents may be stale. In this case, the policy information may be useful
429 because it can provide a rationale as to why the data resource should be considered valid over
430 that refresh cycle. The analysis indicates that knowing data is current is a significant concern and
431 it is likely that this set of elements will gain better definition through operational testing.

432

433

434 **Other metadata bins in this category:**

435

436 Currently, resources tend to fall under either data resources or processing resources. This
437 metadata bin would be expanded if additional resource type is defined.

438

439 **Takeaway**

440

441 The metadata described in this section is the culmination of the groundwork laid from the building
442 blocks on. More atomic building blocks are repeatedly reused, enabling an immediate degree of
443 interoperability even if a user would not understand unique metadata added at a higher level. For
444 example, the Responsible Party instances reuse the structure defined by Person/Organization
445 and thus it should be relatively easy to find a contact point to explain other concepts. The
446 Access/Invocation metadata provides a common description of a resource access point, where
447 the description supports a range of activities that include whether access is allowed or useful at
448 all.

449

450 As noted, several of the metadata bins support instances that will change during the associated
451 entity's life cycle, and such changes will be made by authorized agents other than the original
452 metadata producer. This implies that Resource metadata, at the top of the pyramid, must also
453 have support for update. In addition, resources are assumed to be long-lived and their function
454 and mission space may evolve over time. A metadata system must support and provide
455 configuration management as the building blocks included and the Upper Level metadata
456 referenced change over time to match the new contexts.

457

458

459 **2.2 Policies and contracts**

460

461 Broadly speaking, a policy represents some form of constraint or condition on the use,
462 deployment or description of an owned entity. Policies are inherently unilateral – any participant
463 may have policies about issues that are important to them. A contract, however, is a policy that
464 has been agreed to.

465 A contract can refer to everything from the detailed description of the service interface to the legal
466 contract entered into when two or more parties use a service. However, the SOA RM focuses on
467 those agreements necessary for a successful interaction with a service.

468 **2.2.1 Service Policy**

469 Abstractly, a policy is an assertion that expresses intent on the part of a participant.

470 Policies apply to many aspects of SOAs: to security, to privacy, manageability, Quality of Service
471 and so on.

472

473 Policy assertions may be, but need not be, written down in a formal machine process-able form.
474 Languages that permit policy assertions also range in expressivity from simple propositional
475 assertions to modal logic rules. However, the SOA RM is neutral to how a policy is represented.

476

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

480

481 Associated with policies is the concept of enforcement. Enforcement is the realization of the
482 policy: an un-enforced policy is simply an abstract logical proposition. However, how a policy is
483 enforced, or even whether a policy is enforced is not a relevant part of the reference model.

484

485 A policy always represents a participant's point of view. For example, a provider of a service may
486 have a policy that all users of the service must be authenticated prior to their access to certain
487 functions. This policy is one that may be enforced by the service provider independently of any
488 agreement from potential users of the service. Similarly, someone's agent may embody a privacy
489 policy independently of any services the agent interacts with.

490 **2.2.2 Service Contract**

491

492 Where a policy represents an assertion from the point of view of a participant, a contract
493 represents an agreement between two or more participants. Like policies, contracts can cover a
494 wide range of aspects of services: quality of service agreements, interface and choreography
495 agreements and commercial agreements. However, the concept of a service contract within the
496 SOA RM applies primarily to the requirements for the successful use and provision of services.

497

498 A contract may be, but need not be, expressed in a machine process-able form. It seems
499 significantly likely that an executed contract will not be in a machine process-able form; especially
500 for commercial agreements. However, languages that can express policies, especially the more
501 powerful variants can often also be used to express machine process-able contracts.

502

503 Each contract may be associated with a life-cycle. This life-cycle has three main phases: a
504 negotiation phase, an active phase and a completion phase.

505

506 While it is possible that a specific negotiation phase precedes an agreement to a contract, often it
507 is more implicit. For example, merely attempting to interact with a service may represent an
508 agreement to follow the prescribed procedures for using the service.

509

510 Often a contract specifies policies that are assumed to be in force during the active phase of the
511 contract. As such, those policies are subject to enforcement in a similar way to unilateral policies.

512

513 Enforcement of an agreement will depend on the nature of the agreement: violating an
514 infrastructure-level agreement is likely to lead to errors and unexpected results. Violating a
515 commercial agreement is likely to lead to loss of service or other legal remedies.

516

517 While there may be many kinds of contract, we envisage three main kinds of contract that may
518 apply in service oriented architectures: the contracts that represent the valid use and provision of
519 services, the contracts that represent the permitted uses of services and the contracts that result
520 from using services.

521 For example, the service description may contain descriptions of the interfaces of a service – the
522 kinds of data entities expected and the names of the operations supported – and may also
523 contain choreographic descriptions of the order of interactions. Such descriptions may range from
524 simple identifiers implying a mutually understood protocol to a complete description of the
525 vocabularies, expected behaviors and so on.

526

527 However, a valid use of a service is not equivalent to a permitted use of the service. For example,
528 one may present a syntactically correct request to a service for withdrawing money from an
529 account. If that request is not accompanied by a suitable authentication, then that request is
530 typically denied – it is not permitted. Many security considerations and quality of service
531 considerations lie in this realm of agreement.

532

533 Often the purpose of interacting with a service is to effect a further agreement. For example, one
534 use of a book-selling service is to cause a book to be purchased and delivered.

535 This kind of contract is an important aspect of the rationale for deploying Service
536 Oriented Architectures; however, such contracts are beyond the scope of this SOA RM.


537

538 **2.3 Semantics**

539 A service represents an action boundary between the
540 infrastructure that the service is deployed over and the business
541 context in which it is deployed. Addressing the semantics of this
542 boundary in the appropriate manner is one of the key challenges
543 to developing large scale reliable systems.

544 The semantics of a service are the shared expectations about the
545 service. In many environments, this cannot be represented in a
546 monolithic fashion: the semantics of services have a natural
547 layering.

548



549 **2.3.1 The layers of a semantics of service**

550 Fundamentally, we expect that all services deployed in a SOA have an intended purpose. That
551 purpose is the linchpin by which we measure the expectations for a service and is the basis of its
552 semantics. The purpose of a service is the highest-level semantic characterization of the service.

553 The requirements for reliably and mechanistically interacting with a service represent a baseline
554 for the semantics of a service. This includes any metadata required to contact the service; but
555 also includes such aspects of a service as message transport, data encoding and so forth.

556 The requirements and expectations for the content of any data interchanged. This corresponds to
557 the data model of SOAs

558 The requirements and expectations for the appropriate sequences of interactions. This may
559 include dependencies relating to the stateful requirements on the entities interacting via the
560 service.
561 The requirements and expectations about the intended effects of any interactions – or sequences
562 of interactions.
563 The policies and contracts that may be relevant to the service interaction. Policies and
564 agreements may apply to all levels of the semantics of a service.
565 In principle, the semantics of a service reflects many aspects of its establishment – from the
566 format and structure of any data communicated between the participants of a service interaction
567 to on the participants to the expected effects of successfully interacting with the service.

568 **2.3.2 Metadata**

569 One of the hallmarks of a Service Oriented Architecture is the degree of documentation
570 associated with it. The purpose of this metadata is to facilitate integration, particularly across
571 ownership domains. By providing descriptions, the task of designing client applications that make
572 use of a service is considerably enhanced.
573 In this spirit, we might also expect that the different semantic aspects of a service outlined above
574 may also be documented. Such documentation may be in machine process-able form – in which
575 case it is commonly referred to as metadata – or it may be in informal written form – in which
576 case it is commonly referred to as documentation.
577 If documented in metadata, a service's semantics has many possible uses: it can be used as a
578 basis of discovery in dynamic systems, it can assist in managing a service, validating and
579 auditing uses of services may also be simplified by rich metadata.
580 However, it is not essential to the concept of SOAs that the semantics of a service be so
581 completely described.

582 **2.3.3 Vocabulary**

583 For successful interactions, the various entities must have a shared understanding of the content
584 of interaction as well as the expected behaviors. This includes the meaning of the symbols and
585 strings used in the communication – the shared vocabulary.
586 In some cases, the shared vocabulary can be as simple as a shared data model schema.
587 However, in the context of the Internet, with applications spanning ownership domain boundaries,
588 we are often forced to deal explicitly with meaning – because we cannot rely on the same
589 understanding of terms when different systems are integrated.

590 **2.3.4 Context**

591 Since words and symbols used in a particular discourse may have multiple possible
592 interpretations, and since different participants may have different terms for the same concepts,
593 providing a basis for selecting the correct interpretations of words and symbols is a key to
594 building reliable systems. Context metadata surrounding a particular discourse helps to establish
595 correct interpretation of actions and data between the participants involved in that discourse.

596 **2.3.5 Autonomy**

597 Autonomy is inherently a relative concept -- one is autonomous from control by something. In the
598 case of a SOA-style system, we expect that a very common situation is one where the providers
599 of services and the consumers of services will often belong in different ownership domains; as
600 such, they will inevitably have certain rights and freedoms not normally applicable to closed
601 systems. In particular, providers of services can refuse service, and consumers of services may
602 arbitrarily abandon connections.

603 In this context, it is advisable that SOA architectures be designed from the perspective that
604 service providers and consumers are autonomous from each other. Such a constraint for service
605 participants leads to a reliability benefit in SOA architectures – they will be inherently more robust
606 and reliable than closed architectures.

607

608

609 **2.3.6 Data/Information Model**

610

611 The goal of SOA Data/Information Model is to specify an abstract interface and data model for
612 exchange of data among SOA entities. Entities in SOA need a standard way to (de)serialize data,
613 extract and/or construct metadata, and infer service semantics.

614

615 At the highest level data in SOA can be classified as private and public. Private data includes the
616 data used by a service. The data model of this data is private to the service implementation. This
617 data model does not provide a mapping or bridge to private data. In an exchange this type of data
618 is carried as payload inside of an exchange data unit.

619

620 Public data includes data that embodies the state, property and parameters of an SOA. Public
621 data should be available at standard interfaces and in standard formats.

622

623 The data model should also define standard mechanisms for services to extract metadata that
624 may be serialized with data.

625

626 **2.4 Discovery, Presence and Availability**

627

628 **2.4.1 Discovery**

629

630 Discovery, in the context of Service Oriented Architecture, is the act of detecting, identifying,
631 understanding and selecting a service within the constraints and boundary conditions of a service
632 fabric. It must immediately be pointed out that it is the service description that is discovered not

633 the service itself. A service can only be consumed once it has been located and the interface
634 proffered by the service description is enjoined within the consuming entity.
635

636 **2.4.2 Structured vs. Unstructured Discovery**

637

638 By and large, unstructured service descriptions do not lend themselves to be understood by a
639 large audience. It is reasonable then for sagacious consumption that service descriptions are in
640 the form discussed in section 2.1.2 in order to render them consumable resulting in larger client
641 uptake. However there is no reason other than the rationale presented here and above that a
642 service description be well structured. It is up to the provider to decide a service description
643 format that best suits his/her needs and those of the intended consumers.

644

645 **2.4.3 Service Fabric Constraints**

646

647 Should a service participate in a fabric where the fabric dictates the structure of service
648 descriptions, the provider must then tailor his service description to the accepted format adopted
649 by the fabric management entity. Although service description constraints probably exert
650 restrictions onto the service provider, the “common model” approach allows for a baseline of
651 known semantics and quicker application incorporation.

652

653 **2.4.4 Discovery Methods**

654

655 In general, there are two primary methods by which an entity can be informed about the existence
656 of another and they are:

657 *Discovery by broadcast*, which is autonomously receiving information about an entity or,

658 *Discovery by detection*, which is seeking information about other entities on one’s own volition
659 either intentionally or accidentally.

660

661 The following diagram illustrates the detection models.

662

663

664

665

666

667

668

669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705




Figure 2-2 - Service Description detection methods

The broadcast method for information (in our case service) discovery is used in a number of technologies and on the Internet today with some success; however there are some known limitations and issues related to this method. For example, if an entity issues an information broadcast, specifically a service description “push”, and the other entity (service consumer) is not in a state capable of receiving it, the information is not captured, cannot be acted upon, and is considered lost.

As previously mentioned, discovery by detection has two variants, accidental and intentional. Accidental detection is haphazard at best and reward less at worst. A consumer looking for a suitable service could search previously known locations and if unsuccessful, could then search other random locations ad infinitum. As well, a consumer could send out a broadcast query fabric wide (if supported by the fabric) and then select from all replies within a certain time frame the service that best meets its need. But, how does the consumer know all replies have been received and hence the choice made the correct one?

Intentional discovery by detection is typically through a known medium. This method of discovery has seen support from the standards bodies and significant investment by the industry. Technologies such as registries and standardized search engines provide for well-defined query semantics simplifying and removing the burden from the service consumer and placing it onto a

706 well-known entity or agent (which may also be a service unto itself). There are several competing
707 standards that satisfy the requirements for intentional discovery by detection each with known
708 issues and limitations.
709

710 **2.4.5 Identification, Understanding and Matching**

711
712 Once a service description has been located by whatever means, it needs to be identified and the
713 contents ultimately understood. For example, should two service descriptions resulting from a
714 service query have the same name, there must exist a way to uniquely identify them within the
715 fabric in order to classify them appropriately and consume them intelligently within the context of
716 the consumer and provider. Unique identifiers can be easily constructed based on UUIDs or URIs
717 or some other method designed by the fabric management entity (should one exist) but it is
718 ultimately up to the fabric owner(s) to decide if a identification method is necessary, then agree
719 upon its design and enforcement policy.

720

721 As part of the syntactic coupling required by service consumption, service resource (and
722 specifically service parameter) matching is an integral part of the overall semantic compatibility
723 model. Understanding of intent is as important as the understanding of the resource
724 requirements. For example, in digging a hole in the ground, it is equally important to understand
725 what a hole in the ground is, as having the shovel by which to dig it.

726

727 The SOA Reference Model illustrates where the semantic component has bearing and influence.
728 Understanding of the service description via a normalized service description template (or some
729 such design artifact) assists in the understanding of the service. If the service description can be
730 parsed and dropped into some object that is easily understood by a human or by a machine
731 within the consumer's context, then the service description has achieved its syntactic and
732 semantic goals.

733

734 **2.4.6 Presence and Availability**

735

736 By definition, availability is the ability of an entity to be utilized or consumed within the context of
737 its environment. Previously we have distinguished between the service description and the
738 service itself, and based on this and the definition of availability, service descriptions and services
739 must be discussed separately.

740

741 The availability of a service description is straightforward. A service description is either present
742 in some form on the fabric or it is not. In other words, a service description can be discovered or
743 not. In the case where a service description is discoverable but in a form that cannot be
744 understood by a consumer, it should be considered non-existent to that consumer.

745

746 A service however has significantly more latitude as to its availability. A service description may
747 indicate, for example, certain hours of operation and hence the service need only be operative
748 during those hours. A service can also be in a failure state unable to respond say due to a
749 hardware fault, but its service description is still available implying the service itself is available
750 when in fact it is not. It is completely up to a derived architecture to specify the operational service
751 model for unavailable services and the runtime availability of service descriptions relating to failed
752 or state based services.

753

754 Furthermore for services, one can extend service availability to the “flexible execution” model
755 whereby a service is always available but only executes or instantiates when a “message” is
756 actually received by the service endpoint to consume it. This model, used by dynamic and agile
757 computing technologies, is becoming more and more prevalent as the industry moves to
758 virtualization of the computing enterprise. Of course, the resources need to be available for
759 service invocations should an enterprise utilize this model, but as with all virtualization
760 techniques, there is a point at which the resource base is exhausted and “requests” must be
761 refused or stored for later execution.

762

763

764

765

766

767

768

769

770

771

DRAFT

772

3 Conformance Guidelines

773

774 The authors of this reference model envision that architects may wish to declare their architecture
775 is conformant with this reference model. In order to be conformant to this reference model, a
776 mapping must be made from each core element of this reference model to components of the
777 conformant architecture.

778

779 The following guidelines must be followed for an architecture to be conformant with the SOA
780 Reference Model:

781

- 782 • All services shall be opaque [see Services and Service Composition]
- 783 • Every service shall have precisely one canonical service description [see The Reference
784 Model]
- 785 • Every service description shall contain at least the following elements [see Service
786 interface]:
 - 787 ○ Data Model [see Data/Information Model]
 - 788 ○ Policy [see Service Policy]
 - 789 ○ Contract [see Service Contract]
- 790 • There shall exist a mechanism to convey awareness of a service to all consumers [see
791 Discovery, Presence and Availability]
- 792 • Every service shall advertise their service description via this mechanism [see Discovery,
793 Presence and Availability]

794 **4 References**

795 **4.1 Normative**

796 [RFC2119] S. Bradner, *Key words for use in RFCs to Indicate Requirement Levels*,
797 <http://www.ietf.org/rfc/rfc2119.txt>, IETF RFC 2119, March 1997.
798

799 **4.2 Non-Normative**

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

DRAFT

802 Appendix A. Glossary

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

806

807 Advertising (or Announcement of Availability)

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

810

811 Agent (requester or provider)

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

813

814 Architecture

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

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

820

821 Authentication

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

824

825 (Service) Consumer

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

827

828 Contract

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

830

831 Data Model

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

834

835 Discovery

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

837

838 Encapsulation
839 The act of hiding internal specifications of an entity from the user of that entity, in such a way that
840 the internal data and methods of the entity can be changed without changing the manner in which
841 the entity is used. What is seen by the user is only an interface, or service.
842
843 Framework
844 A set of assumptions, concepts, values, and practices that constitutes a way of viewing the
845 current environment.
846
847 Interface
848 A named set of operations that characterize the behavior of an entity.
849
850 Mediation
851 The transformation, routing, validation and processing of messages.
852
853 Message
854 A serialized set of data that is used to convey a request or response from one party to another.
855
856 Metadata
857 A set of properties of a given entity which are intended to describe and/or indicate the nature and
858 purpose of the entity and/or its relationship with others.
859
860 Negotiation
861 A process that seeks to establish an acceptable basis for a contract between agents for the
862 provision of a service.
863
864 Ontology
865 Represents an agreement within a specific environment of the meanings to be associated with
866 different concepts and their relations to each other.
867
868 Opaqueness
869 The extent to which an agent is able to interact successfully with a service without detecting how
870 the service is implemented.
871
872 Policy
873 A statement of obligations, constraints or other conditions of use of a given service. When a
874 specific set of entities accept such a policy, a contract is usually established.
875

876 Reference Model

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

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

884

885 (Service) Requester or provider

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

887

888 Security

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

893

894 Semantics

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

897

898 Service

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

901

902 Service description

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

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

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

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

908 and thus determine whether the service meets the expectations and requirements of the
909 consumer. Acceptance of the service description by a consumer does not of itself imply a contract
910 to use the service.

911

912 Service Oriented Architecture (SOA)

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

915 rapid and economical implementation in response to new requirements thus ensuring that
916 services respond to perceived user needs.
917 SOA uses the object-oriented principle of encapsulation in which entities are accessible only
918 through interfaces and where those entities are connected by well-defined interface agreements
919 or contracts.
920

DRAFT

921

Appendix B. Use Cases and Examples (Non-Normative)

922

DRAFT

1 Introduction

923

924

925 This section is non-normative. Employing use cases for increasingly complex scenarios, it
926 explores the requirements for developing Service Oriented Architecture specifications using the
927 SOA-RM.

928

929 Three scenarios of an SOA were considered - a simple, an intermediate and a complex, multi-
930 service example. In the simple case there is only one service, which receives and satisfies a
931 request. The intermediate case expands the simple case to use multiple services in one entity.
932 In the complex case, the scenario encompasses multiple services located in different entities,
933 both interior and exterior to the prime service, which directly receives the request from a
934 consumer. As we move towards more complexity, each service implicitly incorporates all the
935 constraints, requirements and solutions of the simpler services.

936

937 The reference model must be sufficient to guide writers of an SOA specification that will satisfy
938 these scenarios. A given specification does not have to satisfy all the scenarios and the
939 reference model must be flexible enough to support specifications that cover only one or two of
940 the scenarios as well as specifications that cover all three scenarios.

DRAFT

941

2 Use-Cases

942

These use cases will be used to test the adequacy and completeness of the reference model for developing specifications for a Service Oriented Architecture for different classes of real-world instances.

944

945

2.1 Simple SOA

946

This use-case describes the simplest type of SOA; a single service, which supports transaction based security, e.g. SSL and a simple policy. No other QoS is required for the simple case. The scenario is a short-lived, atomic transaction. Consequently, it can be considered to be an ACID transaction, i.e. one that is Atomic, Consistent, Isolated and Durable. As long as a transaction can be completed as an ACID transaction it greatly simplifies the resulting Service Oriented Architecture. There is no need for complex compensation as the pending transaction can be frozen and simply rolled back in the case of failure. Further, the single service eliminates the need for correlating the activity of multiple services. Of course, the service may have a potentially large number of applications to complete its work, but that is out of scope for an SOA specification.

956

957


Figure 1 shows the flow for the simple SOA case. When the service receives a request, it accesses its policy to determine if the request can be satisfied using the simple scenario. If this is true it retrieves its security requirements from the policy. Then the service checks whether the request meets its security requirements. Other situations for refusing the request may exist, for instance, where the provider is unable to handle a request in the time that the consumer requires.

958

959

960

961



962

963

Figure 2-1

964

965

If the consumer requirements are met, the service processes the request and returns the results; otherwise the service performs its fault activity.

967

968 Since the root service may receive multiple requests, it must have some means of identifying the
969 different requests. The identification, which is sent to the requester in the reply, must be
970 understandable to the requester. The service should have some means to advertise its service
971 so the requester can discover it. Note that once the requester discovers a service for a particular
972 activity it may preserve this information and have no further need for the discovery service unless
973 circumstances change.


974

975 Note that the consumer may not wish to disclose certain information in an initial service request
976 until it can be certain that its request will be acceptable to the provider, by virtue of the fact that it
977 conforms to the provider's policy. Thus, there may be multiple requestor messages to complete a
978 single request.

979

980 When a fault occurs, see figure 1.1, the service performs a rollback and reports the fault to the
981 requester. Since the transaction is short-lived and contained within one service, the simple
982 scenario must be designed to return to its pre-existing state without any complex activities.

983



984

985 *Figure 2-2*

986

987 **2.2 Intermediate SOA**

988 An intermediate SOA scenario differs from a simple SOA in that it is composed of multiple
989 services, which are located in the same domain and under the control of one entity.

990

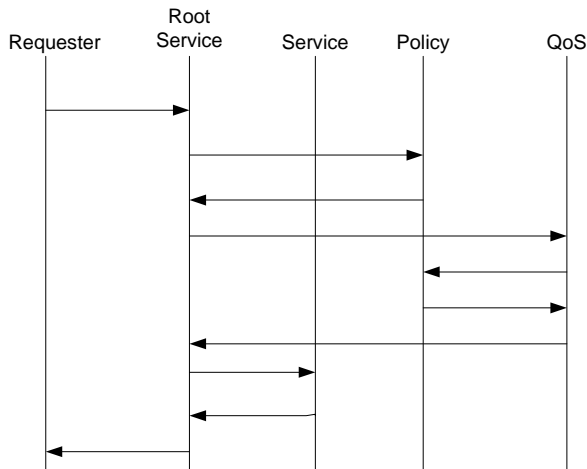
991 The intermediate server provider satisfies requests by using multiple services behind a single
992 service façade, the root service, exposed to the requester. If the service request conforms to the
993 provider's policy for requests, the provider accepts the request. Otherwise, it returns a fault.

994

995 **Requester** **Service** **Rollback** **Fault**
996 As opposed to the simple SOA scenario, the root service uses additional services to complete an
997 activity. The root service controls the activity of the secondary services since they are all part of a
single entity. Since there are multiple services, there is a need for coordination of the services.

998 (Coordination has the meaning here of controlling the flow of the process and is not meant to infer
999 any existing technology.) Some of the services in this scenario may be long-lived and thus do not
1000 have the ACID properties of the simple use-case. In addition, some of the services may require
1001 intermediate information from other services to complete their activity. In the later case the root
1002 service should request the intermediate information from the supplying service and deliver it to
1003 the service or services that need the information.

1004



1005

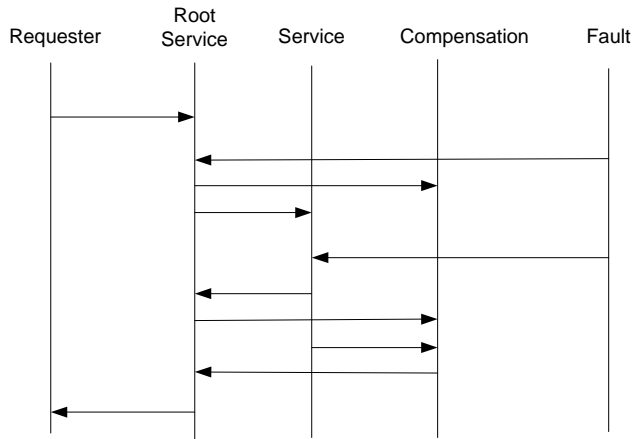
1006 *Figure 2-3*

1007

1008 An SOA specification aimed at the intermediate scenario must be able to handle additional
1009 Quality of Service requirements beyond that of the simple service. These would include
1010 management of the secondary services and long running transactions.

1011

1012 In the situation where the service cannot complete its activity it must supply a compensating
1013 activity and report the fault to the root service. In turn the root service must be able to
1014 compensate for the fault or transmit the fault to all the secondary services supporting the activity
1015 and report the failure to the requestor. Each of the secondary services must perform their
1016 compensation activity.



1017

1018 *Figure 2-4*

1019

1020 **2.3 Complex SOA**

1021 A complex SOA scenario is one in which a requester submits a service request to a service
 1022 provider, which requires child services to complete their requested activity. As opposed to the
 1023 intermediate scenario, the child services used by the complex service may be located in the root
 1024 service entity as well as other locations, which are controlled by entities distinct from the root
 1025 entity. The secondary layer of services may use other services, again from other entities and
 1026 which are independent of the secondary layer of entities forming a tree of arbitrary depth.

1027

1028 In this scenario, there may be short run services, which may be treated as atomic transactions
 1029 and long running services, which may require hours or days to complete their tasks.

1030

1031 A complex SOA architecture has the following characteristics and capabilities:

- 1032 • A hierarchy of services, i.e. the primary service must be able to control and use
- 1033 secondary services, which may in turn control and use tertiary services, etc.
- 1034 • The ability of the services in each layer must be able to support a two phase commit
- 1035 and/or report and compensate if they fail or are informed of a failure elsewhere in the
- 1036 system.
- 1037 • A child service may have more than one parent.
- 1038 • A child service may have no direct knowledge beyond their parent service(s).

1039

1040 The primary or root service must be able to interpret incoming requests and determine:

- 1041 • What tasks are required to satisfy the request
- 1042 • What additional services are required to complete all the tasks

1043

1044 Once the primary service determines the workflow for the activity, it then has to:

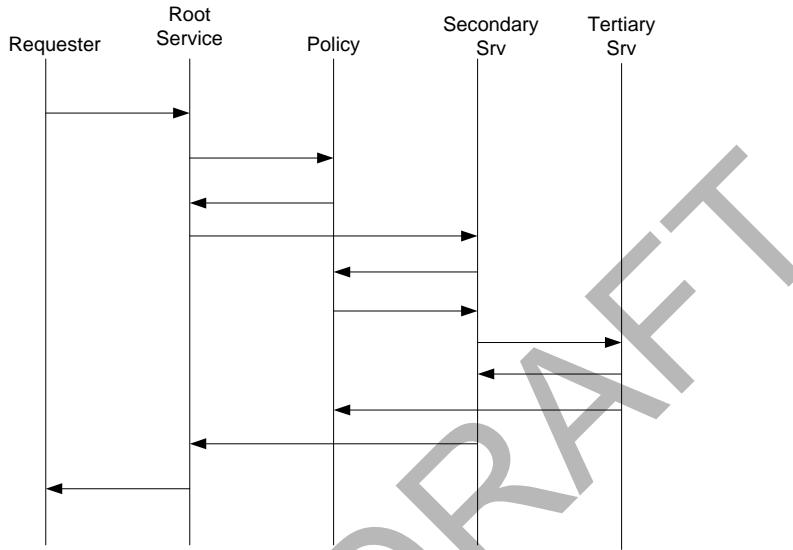
- 1045 • Discover and alert each of the required child services to prepare to accomplish their tasks
- 1046 • Pass any required intermediate results to services that require the intermediate information. This information may pass through multiple layers of the hierarchy
- 1047
- 1048 • Once all the services have reported they are prepared the root service instructs them to complete their activity.
- 1049

1050

1051 All services must possess the ability to:

- 1052 • Advertise their services so that others may use them
- 1053 • Compensate for an abort caused by any faults that may occur in any of the services including their own
- 1054
- 1055 • Understand and be able to act on instructions from their parent service(s).

1056



1057

1058 *Figure 2-5*

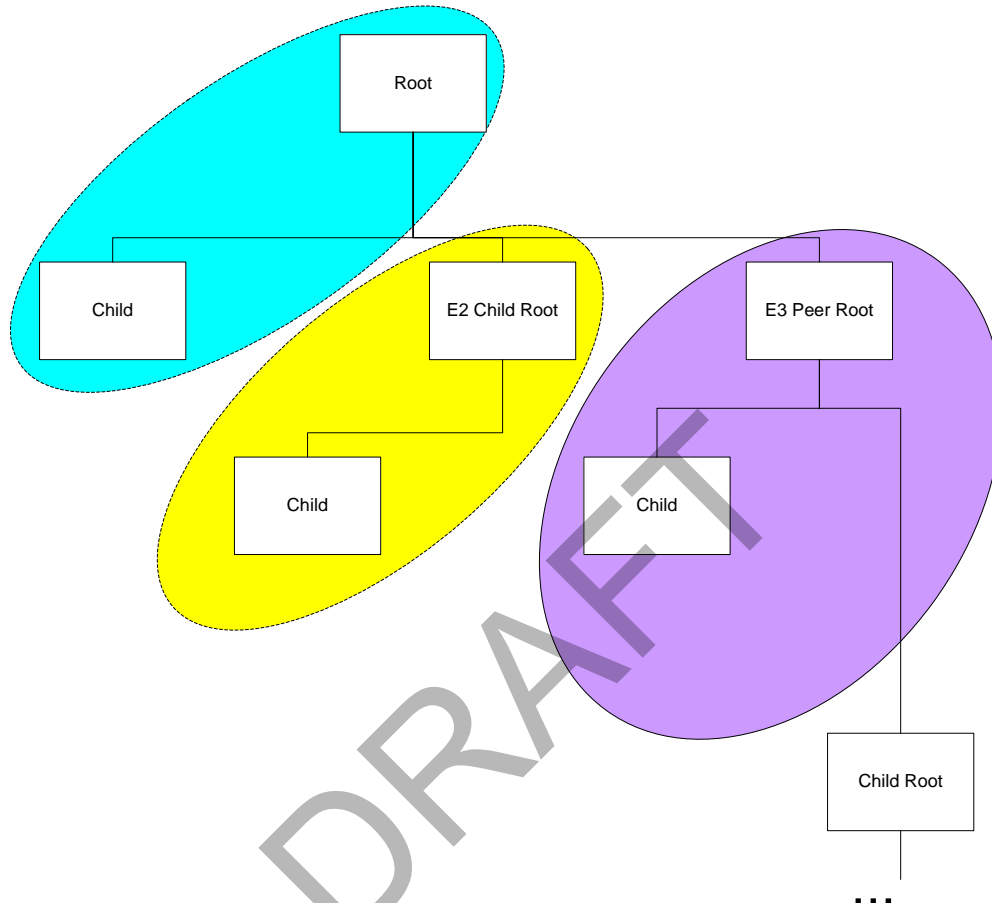
1059

1060 Figure 2-5 shows the activities for a complex service hierarchy. In this scenario the primary
 1061 service depends on a number of layers of services to complete its activity. Lower levels of
 1062 service may depend on further lower levels of services to complete their activities. Each service
 1063 must satisfy the requirements of a simple service and possibly an intermediate service.

1064

1065 The relationship between different layers of services may be peer-to-peer or parent-child. Figure
 1066 2-6 shows an example of the hierarchal structure of the complex scenario. The colored ellipses
 1067 represent the different entities. The blue ellipse represents the primary entity. The primary entity
 1068 contains the root service for the activity, which has one or more children contained within its
 1069 entity. The yellow ellipse represents another entity, E2, whose root entity has a parent child
 1070 relationship with the primary root service. The purple ellipse, E3, represents an entity that has a

1071 peer-to-peer relationship with the primary root. Entities E2 and E3 may have further peer-to-peer
1072 or parent-child relations with other entities.
1073



1074
1075 *Figure 2-6*

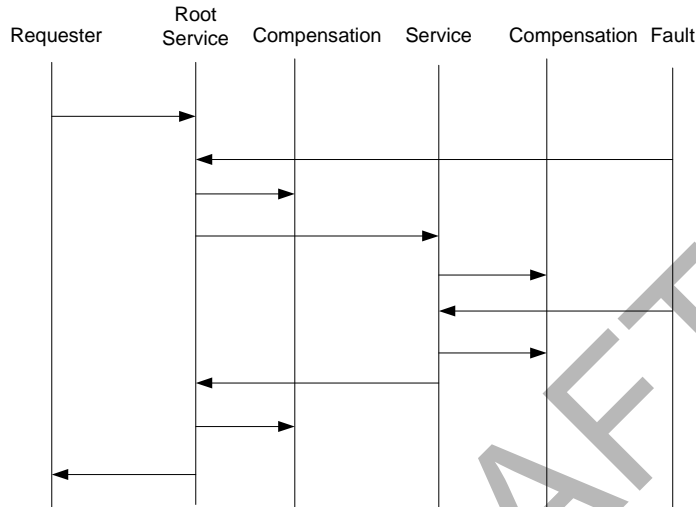
1076
1077 In the peer-to-peer case, each set of services is under control of a different entity and may have
1078 policy controlled by that entity. This will require policy and control negotiation between the
1079 different entities. In the parent-child case the parent entity dictates the policy and control
1080 structure of the child entities. Note that as in Figure 2-6 there may be a mixture of the two types
1081 of relationship to satisfy of a given SOA process.
1082

1083 Any service in the hierarchy may produce a non-recoverable fault, which must be passed up the
1084 hierarchy. When it reaches the root service it is the duty of the root service to compensate for the
1085 fault or report the failure to the requestor.

1086

1087 Figure 2-7 shows the fault scenario for the complex use-case. The major difference from the
1088 intermediate case is that each of the disparate entities has its own compensation and the
1089 disparate entities relay the results of their compensation to the root service. The root service is
1090 responsible for send the fault related to the activity to other secondary services and to the
1091 requestor. Each entity root service is responsible for sending the fault to its children.

1092



1093

1094 *Figure 2-7*

DRAFT

1095

Appendix C. Metadata in the context of a SOA

1096

1097 Metadata is often described as a critical element that will support and enable a service-oriented
1098 architecture. To accomplish the many functions for which it has been associated, metadata for a
1099 SOA must go beyond being the data model in a database or the information included before a
1100 table in a data file to identify the variables represented by the values in the rows and columns.
1101 For example, metadata has been discussed in terms of the following capabilities:

1102 Consumers must be able to search for resources without knowing the details, such as specific
1103 APIs, of the resource beforehand. This implies that the description of the resource must be
1104 expressed in a universally accessible format and, though it will be associated with the resource,
1105 the description will be external to the resource so it can be accessed without reading or otherwise
1106 invoking the resource itself.

1107 The external description must contain sufficient detail so the consumer can decide if the resource
1108 will satisfy the current need.

1109 If the resource is appropriate, the consumer must be able to access the resource content or
1110 invoke the resource processing without knowing the APIs or other details of the resource.

1111 If the consumer attempts to access the resource, sufficient information must be available about
1112 the consumer so that the provider or an agent acting for the provider can determine if the access
1113 is authorized.

1114 The producer and consumer must share a common format for the description and must also
1115 agree on how to interpret the description content. This may be accomplished by indicating a
1116 common vocabulary or distinct vocabularies for which services exist to mediate a translation.

1117

1118 To accomplish this, the traditional definition of metadata must be expanded. In the SOA context,
1119 we will define metadata to be *a subset of the data related to an entity that provides some critical
1120 descriptive information which is useful in some context for identifying, using, or otherwise
1121 interacting with the entity.* It provides a set of descriptive properties which serves one or more of
1122 the following functions

1123 uniquely characterizes an entity and for which values associated with the descriptive properties
1124 allow a user (human or machine) to discriminate between one entity and another,

1125 describes how the entity and its contents can be accessed (both procedurally and the terms of
1126 access) in either a read or write mode or executed if the entity comprises processing instructions,

1127 contains pointers to information not explicitly part of a given metadata set but which is required as
1128 processing or control inputs by other applications or services.

1129

1130 Metadata often includes what the entity is, where it is located, and how to make use of it. It may
1131 describe entity properties such as format, structure/organization, context, business rules, or any
1132 other chosen elements of its integral or associated data or capabilities. It may include the calling
1133 argument to methods, invocation of services, or similar executable commands that act on the
1134 content of an instance of the entity, including accessing it from its native storage format.

1135

1136 Examples of metadata

1137

1138 Example 1 – metadata for a book

1139 Consider the ways in which metadata for a book may be defined and used for different contexts.

1140 For a librarian, the Library of Congress classification number is likely an important metadata
1141 element.

1142 Conversely, for a bookseller, the classification number is not likely to be as important but the
1143 current sales price would be (while this price may not be of interest to the librarian).

1144 The text in the book is unlikely to be identified as metadata, but specific quotes from the book
1145 may be metadata for someone advertising the book.

1146

1147 Example 2 – getting the weather

1148 Consider a user looking for meteorological data. Metadata associated with a data resource that
1149 could support this includes

1150 general document metadata with the name of the data resource and the geographic locations
1151 from where it can be accessed; metadata specific to the function of the data resource, such as
1152 the date, time, and location where the data was collected,

1153 access control restrictions which must be satisfied (or possibly licensing terms if it is a
1154 commercial source) and a pointer to the service interface (e.g. WSDL) to retrieve the data,

1155 a pointer to pedigree information describing the quality of the data as evaluated based on how the
1156 data was collected and processed and the accuracy of the measurements.

1157

1158 The request for the meteorological data may generate a log file detailing the services invoked and
1159 resources used to satisfy the request, and the log file could be archived using a network storage
1160 service. Associated with the stored log would be metadata containing a log ID, the date of the
1161 request, and the identity of the requester. Note, in this example, the log file itself is not
1162 considered metadata but information describing the log file is. A pointer to the log metadata
1163 would be returned with the requested data so the requester would both know how the request
1164 was fulfilled and be able to point to the log as a repeatable means to satisfy a similar request in
1165 the future.

1166

1167 As noted in both the book example in the Introduction and the weather example in the previous
1168 section, what constitutes the appropriate metadata set depends on the context of the user and
1169 the current needs to be satisfied. Thus, it is less important to have defined the perfect metadata
1170 set than to ensure that the combined metadata available can provide or support access to the
1171 critical information at the critical time.

1172

1173

1174

1175 **Appendix D. Acknowledgments**

1176 The following individuals were members of the committee during the development of this
1177 specification:

1178 [TODO: insert cte. Members]

1179

DRAFT

1180

Appendix E. Notices

1181 OASIS takes no position regarding the validity or scope of any intellectual property or other rights
1182 that might be claimed to pertain to the implementation or use of the technology described in this
1183 document or the extent to which any license under such rights might or might not be available;
1184 neither does it represent that it has made any effort to identify any such rights. Information on
1185 OASIS's procedures with respect to rights in OASIS specifications can be found at the OASIS
1186 website. Copies of claims of rights made available for publication and any assurances of licenses
1187 to be made available, or the result of an attempt made to obtain a general license or permission
1188 for the use of such proprietary rights by implementers or users of this specification, can be
1189 obtained from the OASIS Executive Director.

1190 OASIS invites any interested party to bring to its attention any copyrights, patents or patent
1191 applications, or other proprietary rights, which may cover technology that may be required to
1192 implement this specification. Please address the information to the OASIS Executive Director.

1193 Copyright © OASIS Open 2005. *All Rights Reserved.*

1194 This document and translations of it may be copied and furnished to others, and derivative works
1195 that comment on or otherwise explain it or assist in its implementation may be prepared, copied,
1196 published and distributed, in whole or in part, without restriction of any kind, provided that the
1197 above copyright notice and this paragraph are included on all such copies and derivative works.
1198 However, this document itself does not be modified in any way, such as by removing the
1199 copyright notice or references to OASIS, except as needed for the purpose of developing OASIS
1200 specifications, in which case the procedures for copyrights defined in the OASIS Intellectual
1201 Property Rights document must be followed, or as required to translate it into languages other
1202 than English.

1203 The limited permissions granted above are perpetual and will not be revoked by OASIS or its
1204 successors or assigns.

1205 This document and the information contained herein is provided on an "AS IS" basis and OASIS
1206 DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO
1207 ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE
1208 ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A
1209 PARTICULAR PURPOSE.