

46 THE SPECIFICATION OR THE PERFORMANCE OR IMPLEMENTATION OF THE
47 CONTENTS THEREOF.

48

49 You may remove these disclaimers from your modified versions of the Specification
50 provided that you effectively disclaim all warranties and liabilities on behalf of all
51 copyright holders in the copies of any such modified versions you distribute.

52

53 The names and trademarks of the Authors may NOT be used in any manner, including
54 advertising or publicity pertaining to the Specification or its contents without specific,
55 written prior permission. Title to copyright in the Specification will at all times remain with
56 the Authors.

57

58 No other rights are granted by implication, estoppel or otherwise.

59

60 **Status of this Document**

61 This document is provided as a submission to the OASIS Web Services Distributed
62 Management Technical Committee and to the Web service community in general.

63

64 **Abstract**

65 As Web services become pervasive and critical to business operations, the task of
66 managing Web services and implementations of the Web services architecture will be
67 imperative to the success of business operations. Web services manageability is defined as
68 a set of capabilities for discovering the existence, availability, health, performance, and
69 usage, as well as the control and configuration of a Web service within the Web services
70 architecture. This implies that Web services can be managed using Web services
71 technologies. The importance of a standardized management model for Web services and
72 the promise of Web services as a management integration technology has driven industry
73 leaders in Web services and management technologies and applications to form a technical
74 committee in OASIS called the Web Services Distributed Management (WSDM) Technical
75 Committee.

76

77 The WSDM TC is chartered to develop the specifications for defining how to use Web
78 services to manage any IT resource – the “WSDM Management Using Web Services”
79 specification. They are also chartered to define the manageability model for a Web service as
80 an IT resource and how to access that model using Web services – the “WSDM
81 Management OF Web Services” specification. WS-Manageability defines the later, the
82 manageability model for a Web service and how to access that model using Web services.
83 WS-Manageability does not provide a “Management Using Web Services” specification.
84 But it does provide insight, illustration, and input into how “Management Using Web
85 Services” should be developed to support Web services and Grid services communities
86 through a concrete use case: Managing Web services.

87

88 **WS-Manageability**

89 WS-Manageability is composed of three documents, a concepts document, a normative
90 specification, and renderings listing. These documents are summarized below:

91

92 **The WS-Manageability – Concepts** document outlines the scope and definitions for the
93 specification. It provides an overview of Web services architecture and implications for
94 management of that architecture. The Concepts document further defines the role of the
95 manager in the Web services architecture and provides practical information on
96 manageability implementation patterns and discovery considerations.
97

98 **The WS-Manageability - Specification** begins by introducing the general concepts of a
99 manageability model in terms of manageability topics, (identification, configuration, state,
100 metrics, and relationships) and the aspects (properties, operations and events) used to
101 define them. These abstract concepts apply to understanding and describing the
102 manageability information and behavior of any IT resource, not just Web services. We use
103 these concepts to organize our approach to Web services manageability.
104

105 The manageability model for Web services endpoint is defined as concrete models in UML
106 using the topics and aspects concepts, without implying any particular implementation or
107 locus of implementation. Appropriate manageability interfaces are defined based on the
108 UML manageability models. While some parts of this model may be useful for modeling
109 the manageability of any IT resource, this specification is focused exclusively on the
110 requirements of Web service endpoints and does not propose a complete generic resource
111 manageability model. The WSDM “Management Using Web Services” specification may
112 incorporate these more generic models.
113

114 **The WS-Manageability – Representation** document provides the interface definitions
115 based on the model as WSDL 1.1 and GWSDL renderings. These definitions are meant to
116 show how the topics and aspects concepts along with concrete models can influence the
117 development of consistent Web services interfaces for accessing the manageability
118 information of Web services. The interfaces illustrate how the manageability model for
119 Web services can be divided into aspects of topics that apply to all manageable resources
120 and aspects of topics that apply only to the manageability of Web service endpoints. The
121 interfaces that may apply to all manageable resources may be incorporated into the
122 “WSDM Management Using Web Services” specification. The details of these WSDL
123 interfaces specific to managing a Web service endpoint may change to remain consistent
124 with the “WSDM Management Using Web Services” specification when it is available.
125

126 This specification depends on Web services as a distributed platform. There are several
127 common functions which are required for management using Web services that are not
128 specific to management and not yet available as standardized Web service technologies.
129 Interim solutions have been proposed in this specification which will be replaced by
130 suitable standards when they become available. It is not expected that the WSDM TC will
131 continue specification of or standardization of the interim solutions.
132

133 This specification defines the manageability model for Web services, and then uses Web
134 services as a ‘worked example’ of how manageability interfaces in general, as well as Web
135 services manageability interfaces, may be rendered based on one approach to management
136 using Web services practices.
137

138
139
140
141
142
143
144

Notational Conventions

The key words “MUST,” “MUST NOT,” “REQUIRED,” “SHALL,” “SHALL NOT,” “SHOULD,” “SHOULD NOT,” “RECOMMENDED,” “MAY,” and “OPTIONAL” are to be interpreted as described in RFC-2119 [RFC 2119]. This is not yet done consistently. This specification uses namespace prefixes throughout; they are listed in Table 1. Note that the choice of any namespace prefix is arbitrary at this time and not semantically significant.

Prefix	Namespace
soap	http://schemas.xmlsoap.org/wsdl/soap/
crm	http://www.gridforum.org/namespaces/2003/17/crm
wsdl	http://www.w3.org/2002/07/wsdl
gwsdl	http://www.gridform.org/namespaces/2003/gridWSDLExtensions
sd	http://www.gridforum.org/namespaces/2003/serviceData
ogsi	http://www.gridforum.org/namespaces/2003/OGSI
http	http://www.w3.org/2002/06/wsdl/http
xsd	http://www.w3.org/2001/XMLSchema
xsi	http://www.w3.org/2001/XMLSchema-instance
cbe	http://www.ibm.com/AC/commonbaseevent1_1

145
146
147
148

Table 1: *Some standard prefixes and namespaces commonly used in this document.*

Unresolved issues with the submission are interspersed in appropriate locations through this specification, are highlighted in yellow.

Table of Contents

151 **1. INTRODUCTION.....6**

152 **2. REQUIREMENTS, DEFINITIONS AND SCOPE.....7**

153 2.1. MANAGEABILITY OF WEB SERVICES 7

154 2.2. SCOPE OF WORK 7

155 2.3. DEFINITIONS 8

156 **3. BASIC WEB SERVICE ARCHITECTURE..... 9**

157 3.1. WEB SERVICES ARCHITECTURE 9

158 3.2. IMPLEMENTATION OF THE WEB SERVICES ARCHITECTURE 10

159 3.3. ANATOMY OF WEB SERVICES IMPLEMENTATIONS 13

160 3.4. WEB SERVICE RELATIONSHIPS..... 17

161 **4. THE ROLE OF MANAGEMENT IN THE WEB SERVICES ARCHITECTURE 19**

162 4.1. ANATOMY OF WEB SERVICE MANAGEABILITY 20

163 4.2. WEB SERVICES MANAGEMENT ARCHITECTURE..... 21

164 4.3. MANAGEMENT PATTERNS 22

165 4.3.1. *Consolidated Interfaces*..... 22

166 4.3.2. *Associated Interfaces*..... 23

167 4.4. MANAGERS AND MANAGEABILITY INTERFACES 24

168 4.4.1. *Implementing the Manageability Interface*..... 24

169 4.4.2. *Multiple Managers Considerations* 24

170 **5. DESCRIPTION AND DISCOVERY OF MANAGEABILITY 27**

171 5.1. DESCRIBING THE MANAGEABILITY 27

172 5.1.1. *Consolidated Interfaces*..... 27

173 5.1.2. *Associated Interfaces*..... 28

174 5.2. DISCOVERY OF MANAGEABILITY 29

175 5.2.1. *Discovery of Manageable Web Services*..... 29

176 *From functional Web Service endpoint to its manageability endpoint* 29

177 *From manageability endpoint to the functional Web Service endpoint*..... 30

178 5.3. DISCOVERY OF MANAGEABILITY CAPABILITIES 30

179 **6. SUMMARY 31**

180 **7. REFERENCES..... 32**

181 1. Introduction

182

183 The emergence of Service-Oriented Architectures (SOA) [1] and industry commitment to
184 Web services and Grid computing [2] requires new management capabilities in order to
185 truly achieve the visions and goals of these initiatives. Because, both Web services and
186 Grid environments, leverage the loosely coupled nature of services, deployed across
187 heterogeneous platforms and ownership domains and can dynamically discover and interact
188 with each other during solution design or execution, the manageability of such
189 environments is more critical and challenging to the business they support. Web services
190 can also cross administrative domains leading to multiple aspects of control requiring a
191 greater emphasis on declarative and real-time management through the use of management
192 policies and service level agreements.

193

194 This submission details a manageability model and architecture that form the basis for the
195 Management of Web Services. It directly addresses the requirements for manageable Web
196 service architecture implementations identified through the work of the W3C Web service
197 architecture group (WSAG) and it's Management Task Force (MTF). This submission also
198 takes into account and leverages the work completed and underway in other management
199 domains, including the DMTF's Common Information Model (CIM) and the Global Grid
200 Forum's Open Grid Services Infrastructure specifications and Common Management
201 Model (CMM) proposals.

202

203 This document is structured as follows: [Section 2](#) introduces the Manageability
204 Requirements defined by the W3C and the Management Task Force, the scope of this work
205 and potential future work. [Section 3](#) gives an overview of the Web service architecture and
206 the role of management within that architecture. [Section 4](#) defines the Management
207 Architecture within which we define manageability of Web services [Section 5](#) introduces
208 the principles of description and discovery of manageable Web services.

209

210 The concepts are used in the "Web Services Manageability – Specification" which defines
211 and details the Web services Manageability Model. The specification provides the
212 representation of the manageability model in WSDL 1.1, and Grid WSDL. It explains how
213 this document's material pertains to the work completed in this area for the Grid,
214 specifically the Common Management Model (CMM).

215

216 **Note:** Even though SOAP is used throughout the document for clarity and examples, Management
217 capabilities can be exposed through other protocols including any protocol capable of supporting the XML
218 infoset.

219

220 **Note:** The techniques used to express the Web services manageability model in schema, WSDL and
221 GWSDL descriptions are techniques chosen based on their availability at the time of this draft. These
222 techniques may need to evolve based on new specifications, best practices and techniques for expressing
223 manageability information in WSDL 1.1, WSDL 1.2, and GWSDL.

224 2. Requirements, Definitions and Scope

225 2.1. Manageability of Web services

226 As Web services become pervasive and critical to business operations, the task of
227 managing Web services and implementations of the Web services architecture will be
228 imperative to the success of business operations. Management of Web services in this
229 case is defined as a set of capabilities for discovering the existence, availability,
230 performance health, usage, control and configuration of resources within the Web
231 services architecture. Like other Web Services, manageable Web services architecture
232 implementations need to be secure, including support for authentication,
233 authorization, integrity, and confidentiality.

234
235 The W3C Web Services Architecture Group has identified a set of requirements
236 specific to the management of Web services architecture and Web services. For
237 convenience, the Manageability Goals from the Requirements document [3] are
238 restated here;
239

D-AC018 The Web Services Architecture must enable the management and provisioning of Web Services.

- AC018.1 Ensure that implementations of the Web Services Architecture are manageable.
 - AR018.1.1 Define a base set of standard metrics for architectural components and their interactions accompanied by guidelines for measurement.
 - AR018.1.2 Define a base set of standard management operations for Web Services Architecture implementations. Management operations includes, but is not limited to, support for configuration control and lifecycle control.
 - AR018.1.3 Define a base set of management events to be issued by the Web Services Architecture implementation.
 - AR018.1.4 Define a standard methodology for accessing management capabilities from the Web Services Architecture implementation.
- AC018.2 Ensure that implementations of the Web Service instances are manageable.
 - AR018.2.1 Define how a web service should expose web service specific metrics, configuration, operations, and events.
 - AR018.2.2 Support the discovery of web service management capabilities.
 - AR018.2.3 Define a standard methodology for accessing management capabilities of a Web Service through the Web Services Architecture implementation.
- AC018.3 Ensure that at least the following types of management aspects are supported: Resource Accounting, Usage Auditing and Tracking, Performance Monitoring, Availability, Configuration, Control, Security Auditing and Administration, and Service Level Agreements.

240

241 2.2. Scope of Work

242 This submission is initially limited in scope to the definition of a sufficient set of
243 management capabilities for the architectural element Web services endpoint ([Web](#)
244 [service definition](#), [endpoint definition](#)). We do expect that the scope of this

245 submission will be expanded to define the manageability of the other elements of the
246 Web service architecture.

247 As a result of these requirements for Manageable Web services the W3C introduced
248 the role of the *Manager* (see [Web Service Architecture and the Management Role](#))
249 which uses the manageability information and capabilities provided by manageable
250 resources. The definition of the *Manager* role, beyond the information model offered
251 to it and how it interacts with the existing roles of the Web services architecture, is
252 outside the scope of the W3C Web Services Architecture draft and therefore also
253 outside the scope of this submission.
254

255 The work also has given considerations to the Grid initiatives that also have similar
256 management requirements for Grid services (themselves Web services). Management
257 of Web services is not purely a technical or physical perspective on management but
258 a service perspective and directly ties into the business objectives of the provider and
259 concerns of the consumer.

260 **2.3. Definitions**

261 These definitions are derived from the W3C definitions surrounding the management
262 architecture.

- 263 ➤ **Web service endpoint** is an association between a Binding and a network
264 address, specified by a URI that may be used to communicate with an instance
265 of a Service.
- 266 ➤ **Manageable** implies that a Web service can be managed by a management
267 system.
- 268 ➤ **Managed** implies that an entity is actively being managed by a management
269 system.
- 270 ➤ **Management capabilities** include identification, configuration, metrics,
271 status, operations, and events offered by manageable entities for management
272 purposes.
- 273 ➤ **Manageability** implies the existence of a sufficient set of management
274 capabilities such that an entity is manageable.
- 275 ➤ **Management** is the utilization of the management capabilities by the
276 management system.
- 277 ➤ **Manageability interface** is that interface through which manageability
278 capabilities are offered.
- 279 ➤ **Management endpoint** is the Web services endpoint offering the
280 manageability interface for management purposes.

281 We define one additional term.

- 282 ➤ **Manageable resources** include the subset of components specific to the Web
283 service architecture (Web services, Web service endpoint Web Service
284 Environment, Clients, Client Environments and Web service Descriptions)

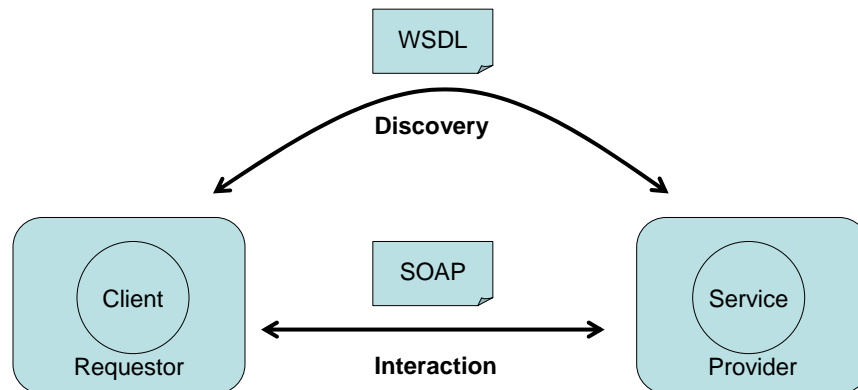
285 and roles defined in a Web service architecture (Discovery Agency,
286 Requestor, Provider, Manager).

287 3. Basic Web Service Architecture

288 3.1. Web Services Architecture

289 W3C defines a Web Service as a software system which functions can be accessed
290 via standard internet protocols by exchanging structural data (XML) messages.

291
292 Although architectural concepts, constraints and relationships are defined in the W3C
293 Web Services Architecture, this document outlines some of the basic principles of the
294 architecture in a more concrete manner to facilitate understanding Web services
295 manageability. We have used the WS-I Basic Profile [4] for Web services, to help
296 explain Web services using XML, WSDL, SOAP, and HTTP technologies. The
297 concepts here still apply when other interaction technologies besides SOAP and
298 HTTP are used.



299

300

301

Figure 1: Web Services Architecture

302

303

304

305

306

307

308

309

310

311

312

313

314

315

316

317

318

Figure 1, represents an abstract architecture with elements such as roles, activities and artifacts. *Roles* are named sets of activities. *Activities* are named sets of procedures (steps, invocations, or interactions). *Artifacts* are named sets of data (information). Every named element has its associated meaning described in detail below.

When actually implemented, an element of a software system may take on one or more roles and therefore be responsible for activities associated with both of those roles.

In most practical cases, a Web service (also referred to as a *service* in this document) is owned by a *provider*. A provider is a role that could represent a legal entity e.g. business or a person owning or willing to provide/offer access to certain functionality via software. A provider has to describe its service in a machine processable manner, that is, in a Web Services Description Language (WSDL) document (artifact). A provider is also responsible for establishing and maintaining an endpoint for a service. An endpoint is an addressable location where data messages could be sent to request execution of some function from the service. The location address, transport

319 details, etc. are usually part of the service description, e.g. a WSDL document, but
 320 can also be defined as a references that is resolved at run-time.

321
 322 A *requestor* is a role that accesses and uses the functionality of the service offered by
 323 the provider. The requestor first discovers the service and understands its
 324 requirements and functions by finding and inspecting its WSDL document. A
 325 requestor is then able to realize its intent by sending messages (artifacts) carrying
 326 structured data (XML) representing the request to the *service endpoint*. Similar reply
 327 messages may then be delivered in response.

328
 329 For interoperability, an endpoint must support a SOAP (WS-Interoperability) binding
 330 for its WSDL interface. The endpoint may also offer additional bindings.

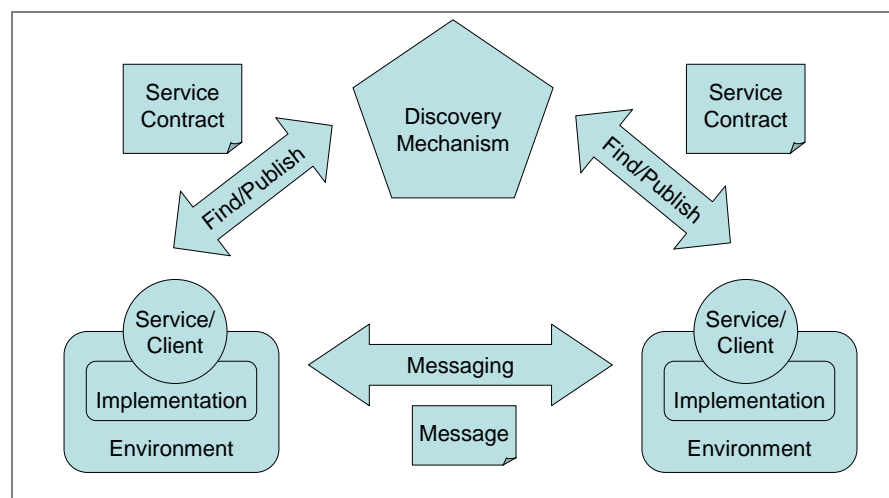
331
 332 The understanding of the provider's intent to offer certain functionality as a service is
 333 encoded in the WSDL document. The requestor's intent to use that functionality is
 334 encoded in the messages sent to the service endpoint.

335
 336 The service and the client in Figure 1 are abstract representations of the entity
 337 realizing the functionality and of the entity intending to use the functionality. An
 338 example of a client could be a Windows application used by a human. An example of
 339 a service could be a piece of software returning a stock quote by accessing
 340 information in a database.

341 **3.2. Implementation of the Web Services Architecture**

342 The abstract architecture may not present enough details for those actually planning
 343 or designing (architecting) Web services-enabled software applications in any
 344 organization. For this reason, this document briefly outlines a deployable
 345 implementation of the Web Services Architecture and the resulting Services Network
 346 (Figure 3).

347



348

349

350

Figure 2: An Implementation of the Web Services Architecture

Note: The depiction above shows the simplest scenario for a Web services architecture but could include and does not preclude support of the concepts of intermediaries and proxies, playing both roles of client and service, between ultimate requester and service provider.

This implementation (Figure 2) builds on the abstract Web Services Architecture and closely resembles other known enterprise software architectures.

When a Web service application is conceived, it has an implementation that is deployed to a suitable environment. An *implementation* is the code that executes, and it is controlled by the *environment* to which it is deployed. Environments can be viewed as purely software (e.g. a platform such as .NET, J2EE) or hardware (e.g. a server) or both. It is likely that many implementations will be deployed to disparate environments but all interact regardless of the vagrancies of the platforms or implementations.

The implementation, often in conjunction with the environment, exposes its functionality as one or more services and, in turn, can consume other services as part of the execution of its functions. For example, an order fulfillment service may use a shipping service to obtain a package tracking number.

Therefore, one implementation could realize abstract service and client roles simultaneously. The service and the client are both abstractions representing exposed and consumed functionality respectively. It is a responsibility of the provider to describe exposed functionality and access details in one or more *Service Contracts*. Service contract is, in many cases, just a WSDL document, but may additionally contain other machine processable markup such as security requirements and policies, e.g. WS-Policy.


The provider is also responsible for publishing service contracts (all pertinent information, technical and business, required for usage of the service) via a chosen *Discovery Mechanisms*. It could use centralized corporate UDDI registry, WSIL files, public search engines, or a peer-to-peer discovery mechanism such as the one found in Open Grid Services Architecture. The presence and use of a discovery mechanism is a key to any sensible implementation of the Web services architecture.

The ability of implementations to dynamically discover and use functionality as Services is one of the great advantages of Web services. Service contracts that define access to abstract functionality could be established independently of its actual implementation or the runtime environment. Properly designed and built Web services applications are loosely-coupled and can be flexibly changed. This enables the logical parts of applications to be dynamically replaced or re-implemented to accommodate new business requirements and changing circumstances.

Service implementations that interact form a logical network of services layered over the infrastructure (e.g. software environments, servers, etc.) and the physical connectivity network. This logical network is also commonly called a Services Network (see Figure 3). As a Services Network grows, its existence and performance

398
399
400

becomes highly relevant to the main line of business, just like other critical IT resources.



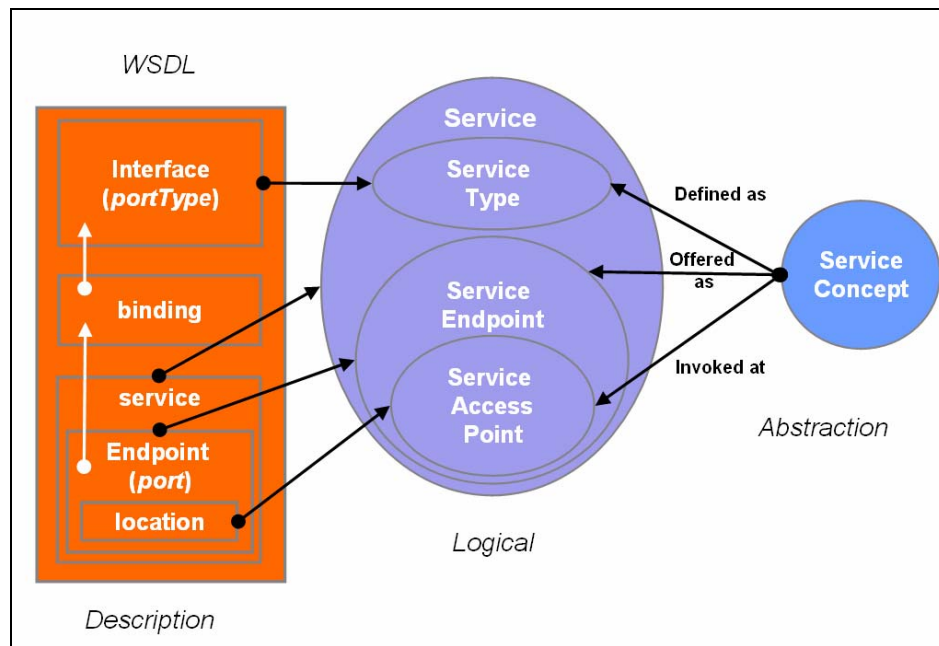
401
402
403

Figure 3: Web Services Network

404 **3.3. Anatomy of Web Services Implementations**

405 Within the Web services architecture the focal elements are the *service* and the *client*.
 406 The abstract idea of a service, once conceived, has to be defined, implemented,
 407 deployed, and offered (not strictly in that order) such that it can be utilized by clients.
 408 A requestor has to recognize the need for a service, discover it, and implement the
 409 client to use the service.

410
 411 Figure 4 depicts the organizing principles and relationships between the elements of
 412 the Web Services Architecture that provide for the materialization of a service
 413 concept into an invocable service. The functional definition of a service or, in other
 414 words, an interface of a service is what determines the service type. A particular
 415 instance of a service type is described as a Web service endpoint and includes the
 416 information required to invoke the service. Definition of a service endpoint includes a
 417 service access point that defines the network location at which the service can be
 418 invoked. In terms of the WSDL, the logical service type is described in the interface,
 419 the service endpoint is an endpoint or port, and the service access point is the address
 420 (location).
 421



422
 423 **Figure 4:** Abstract anatomy of a Service Definition
 424

425 The process for defining, implementing, offering, and potentially publishing a service
 426 can be seen in more detail below (Figure 5). The initial business concept of a service
 427 materializes in the logical and physical architecture as shown. Some elements are
 428 logical in terms of the service being offered, and some are physical in terms of the
 429 implementation. These elements may well be the responsibility of different parties.
 430

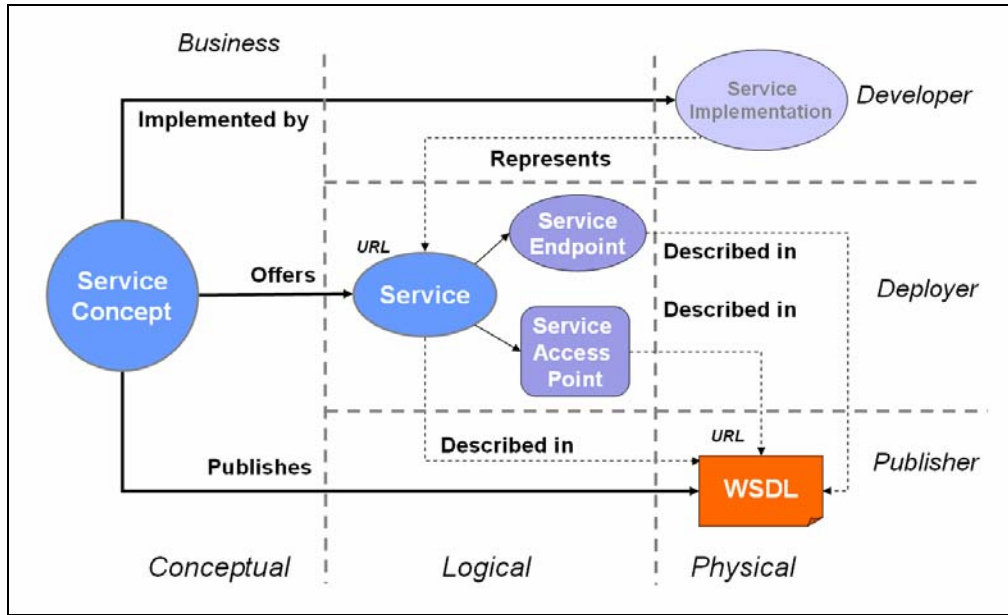


Figure 5: Anatomy of a Service Lifecycle

431
432
433
434
435
436
437

The same perspective can be used to look at the anatomy of a client. Figure 6 shows how a business need for a service materializes by discovering WSDL, implementing a client (requesting application), and using the client to interact with the service.

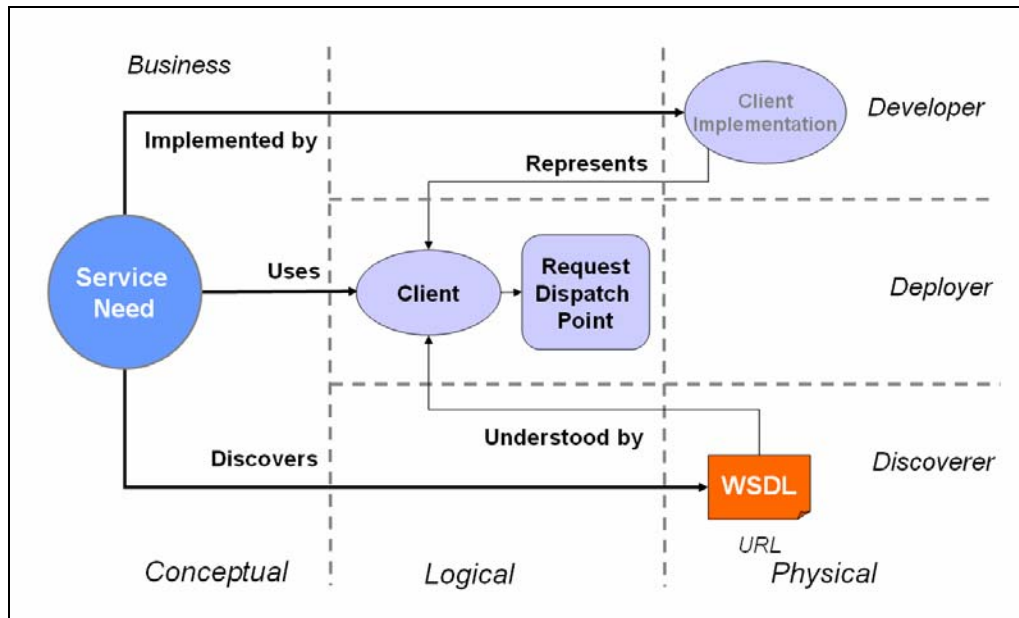
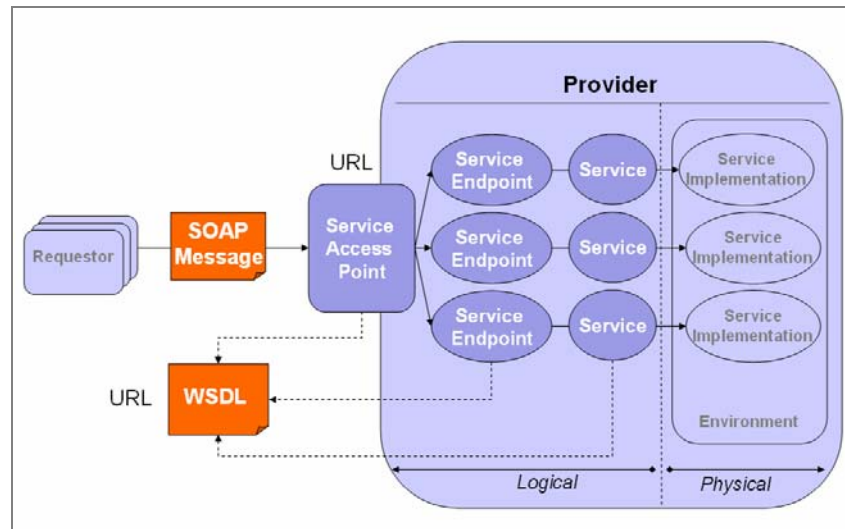


Figure 6: Anatomy of a Client Lifecycle

438
439
440

441 The anatomy of the service and client can be viewed even more concretely by looking
 442 at the implementation and deployment of both elements. Figure 7 depicts the anatomy
 443 of a service within the provider.
 444



445
 446 **Figure 7:** Provider perspective on services
 447

448 Each *service* has an implementation (*service implementation*). It is the code or logic
 449 that implements the functionality offered by the service. It is exposed as a Web
 450 service through a service endpoint or potentially multiple service endpoints. Each
 451 endpoint is essentially a listener for protocol specific messages at the given URL (via
 452 a *service access point*). The URL may provide access to one endpoint or many
 453 endpoints.

454
 455 The *service* itself is the logical element that represents the offered functionality,
 456 which is implemented by the service implementation.

457
 458 The *service endpoint* is the representation of the access mechanisms to the service. A
 459 service endpoint may be supported by several service implementations for scalability,
 460 differentiated service, or SLA enforcement.

461
 462 A *service access point* is the logical element that represents the URL at which
 463 endpoints receive messages and to which requestors send messages. A URL may be
 464 shared by multiple endpoints. In such sharing cases, message dispatch happens based
 465 on the content of the messages and their context. For example, the header of the
 466 message can identify the exact recipient of the routed message.

467
 468 The *environment* hosts or executes service implementations and provides for service
 469 access points. It provides support for a stack of protocols and infrastructure involved
 470 in offering a Web service. For example, it may include HTTP over TCP/IP listeners,
 471 SOAP processor, XML marshalling/de-marshalling, etc. The Environment can be
 472 perceived as a software platform, a hardware server, or an entire data center.
 473

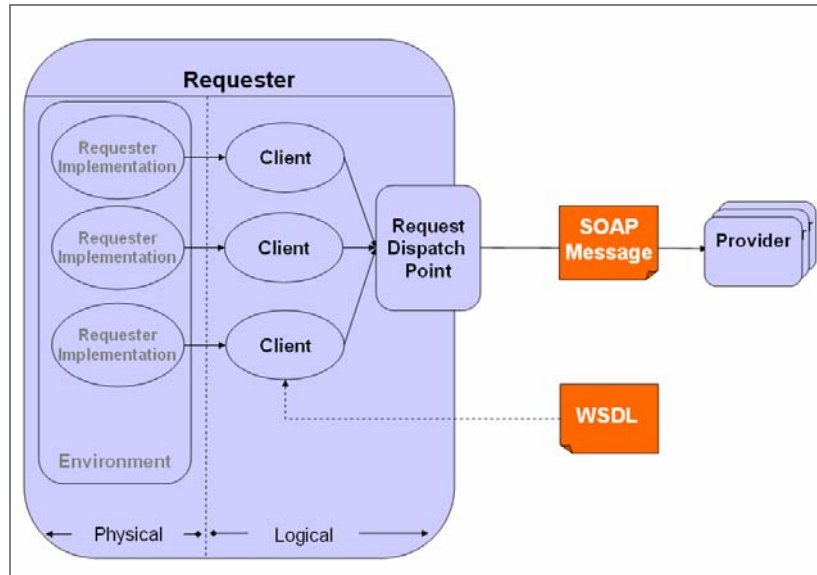


Figure 8: Requester perspective on Clients

474
 475
 476
 477
 478
 479
 480
 481
 482
 483
 484
 485
 486
 487
 488
 489
 490
 491
 492
 493
 494

In Figure 8, the anatomy of a *client* is portrayed within the requester role of the Web services architecture. It also shows the *requester implementation* and its *environment*.

A *requestor implementation* utilizes the formal description of a Web service which it reads from the WSDL document.

The client represents the capabilities natively to the requester implementation such that the Web service can be invoked. For example, a requester implementation may be a Windows GUI application and the client is the part that encapsulates the particular knowledge about a given service, such as how to marshal parameters into XML representation and how to pack a SOAP message properly to the service.

The environment hosts and executes requester implementations and acts as the *dispatch point*. The dispatch point sends messages assembled by the client to one or more providers of the service. It then receives the response from the service. The client further processes it into a form that the requester implementation expects (e.g. a set of Java objects).

495 **3.4. Web Service Relationships**

496

497

498

499

500

501

502

503

504

505

506

507

508

509

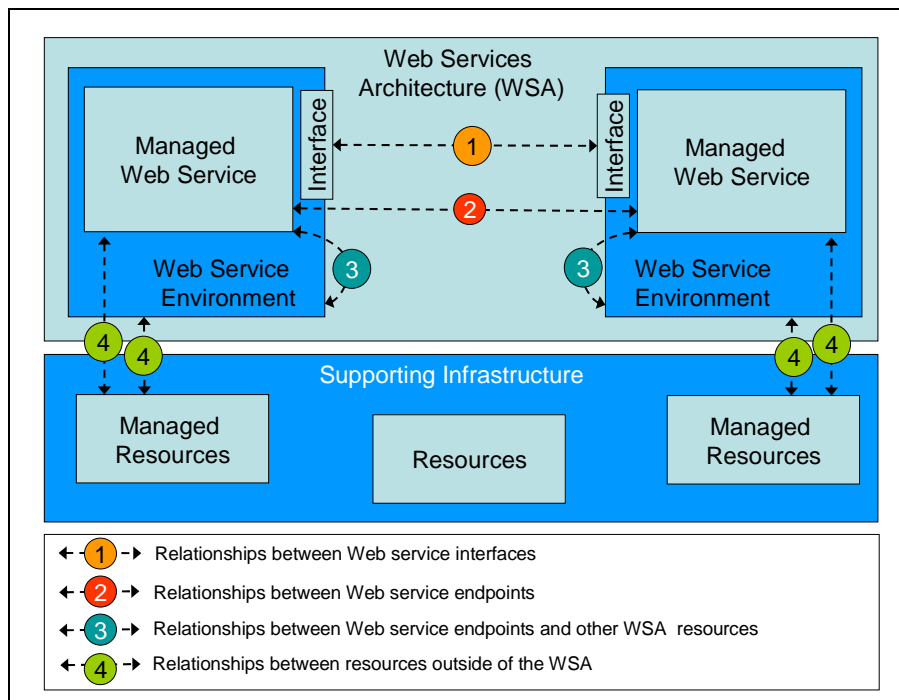
510

511

Intrinsic relationships exist between elements of the architecture (e.g. services, their implementations and the environments they are deployed to), and between Web services themselves. The latter represents the business/functional semantics of the services rather than the way services are realized in the architecture.

Web services interact with each other to offer new composite Web services, and the interactions form the logical networks of services.

The relationships, important to managing Web services, can be summarized into four distinct categories: 1) relationships between Web service interfaces, 2) relationships between Web service endpoints, 3) relationships between Web service endpoints and other architectural elements specified in the Web services architecture, and 4) relationships between Web service endpoints and other manageable resources outside of the WSA.



512

513

514

515

516

517

518

519

520

521

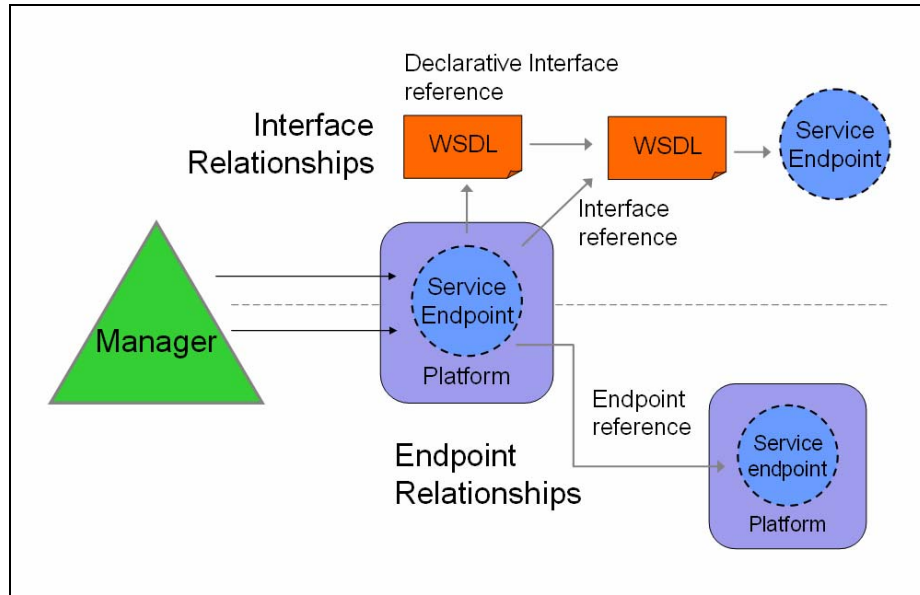
522

Figure 9: Web service Relationship Categories

The relationships defined for a Web service in this document fall into the first two groups – the relationships between Web service interfaces and Web service endpoints. The continued work surrounding the management of the Web services architecture and the underlying infrastructure will no doubt lead to additional relationships currently outside the scope of this document.

523
 524
 525
 526
 527
 528
 529
 530
 531

Web service relationships are declared in two ways - at an interface level or at a service endpoint level. Declaring relationships at the interface level allows for visibility into the relationships even where the actual Web service endpoint is resolved dynamically at runtime. Interface level relationships can be defined declaratively in the WSDL document describing the Web service, or they can be made available by the Web service endpoint itself via manageability interfaces. Service endpoint relationships are only made available by the Web service endpoint via manageability interfaces.



532
 533

4. The Role of Management in the Web Services Architecture

To help represent manageability of the Web Services Architecture, a *manager* role was introduced (Figure 11) by the W3C Web Services Architecture Working Group. A manager is responsible for management of the manageable resources in the architecture (see [definitions](#)).

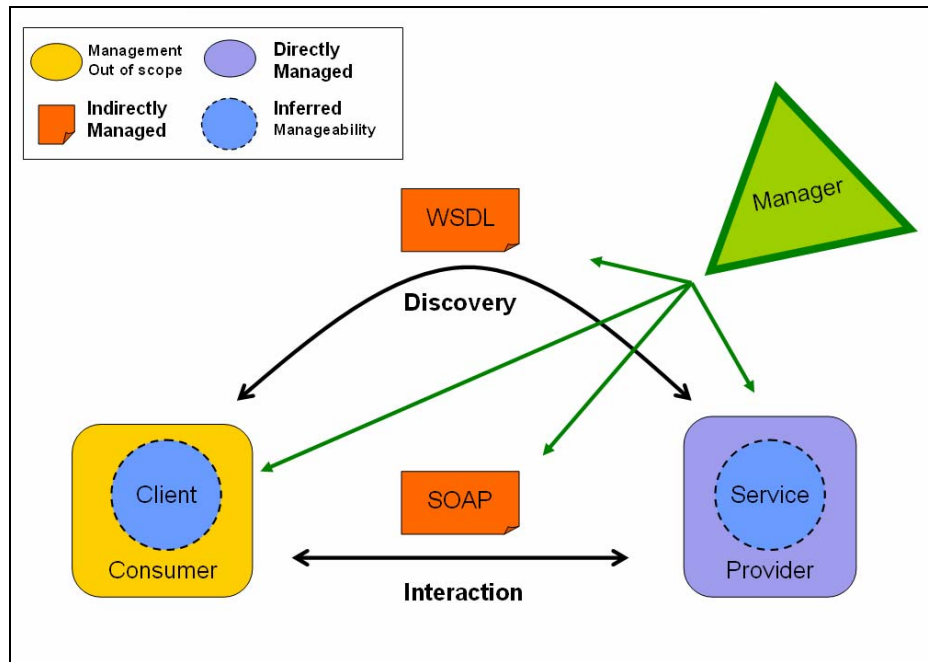
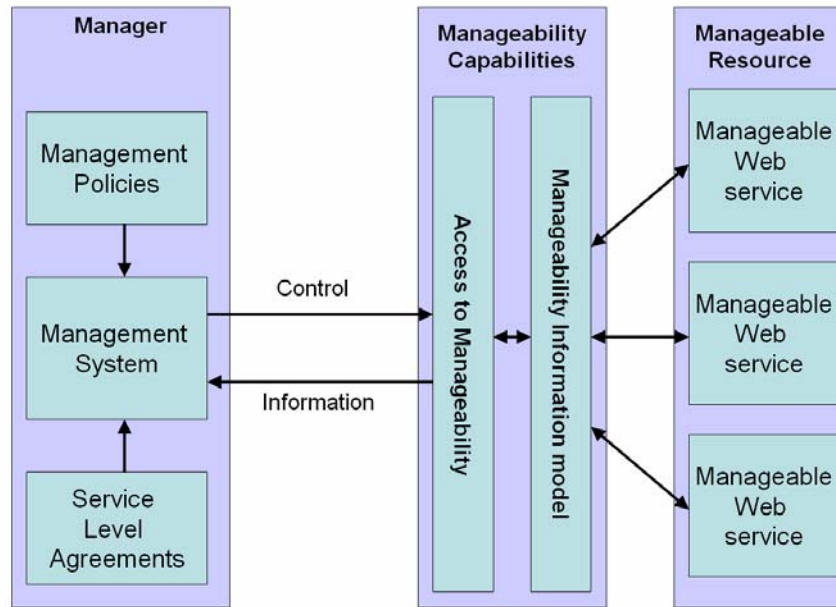


Figure 11: The Manager Role

Various resources of the architecture may have different *manageability capabilities*, and the manager is essentially the “user” of those capabilities (see Figure 12). A manageability capability could be something as simple as enabling/disabling certain function of an element. For example, it could stop acceptance of all requests by a Web service. Manageability capabilities are realized through a *manageability information model* that represents the manageability of resources and manageability information concerning those resources (state, configuration, relationships etc). The model can be accessed by managers which may retrieve information and exert control. For example, semantics of the manageability of a certain resource can be expressed in a UML model, and access to those specific manageability semantics may be described in WSDL interfaces and XML Schema data types. Managers adhering to both standardized access and semantics may manage those resources in an interoperable manner.



557
558
559
Figure 12: Web service Manageability Capabilities

560 Some elements of the architecture are *directly managed*. This means that a manager can
561 exercise manageability capabilities offered by that element directly. For example, a
562 manageability interface may be exposed by the element itself. In the Web services
563 architecture, a service endpoint is an element that can be directly managed.
564

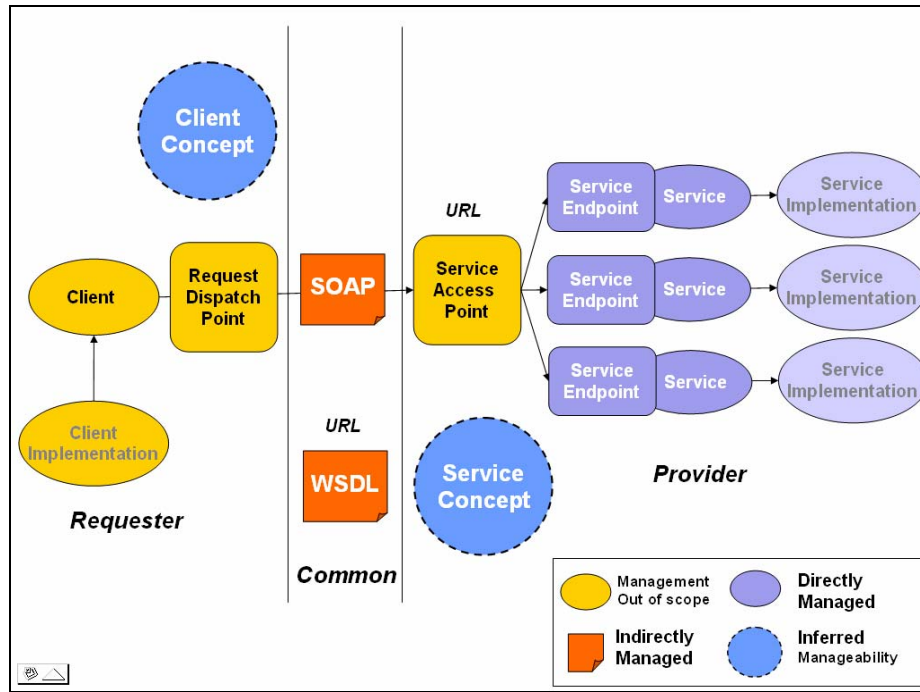
565 The manager is also aware of other elements of the architecture that are indirectly
566 managed. *Indirect management* means the manager can request or subscribe to the
567 information about those elements but doesn't directly access the element for the
568 information. For example, SOAP messages can be recognized, counted, and related to
569 the managed Web service, yet a manager cannot invoke an operation on a SOAP
570 message.
571

572 There is also information important to manageability beyond that provided directly or
573 indirectly by elements of the architecture. This manageability information is *inferred*
574 through analysis of the information provided. For example, a SOAP message header
575 can contain a digital certificate of the client's identity. A manager would have to extract
576 that information and translate into the client's identification that could be used for
577 management purposes later on (e.g. counting SOAP messages sent to a particular Web
578 service by a particular client).

579 **4.1. Anatomy of Web Service Manageability**

580 The anatomy of a Web service discussed in [section 3.3](#) can be categorized according
581 to the type of manageability available. Service endpoints (following through to
582 services) are directly manageable by any actual implementation that assumes the
583 manager role (e.g. a management system). The SOAP messages and WSDL
584 documents are indirectly manageable. Manageability of client and service concepts is
585 inferred.

586



587

588

Figure 12: Anatomy of the Web Service Manageability

589 **4.2. Web Services Management Architecture**

590

591

592

593

594

595

596

597

598

599

600

601

602

603

604

605

A Web service is manageable if its manageability capabilities are exposed via standard manageability interfaces. The manageability interfaces must be accessible in a way compliant with the Web Services Architecture. In other words, manageability interfaces are exactly similar to the functional interfaces exposed by a Web service, only that they bear management related semantics and are used by the manager and not by the client. Figure 13 illustrates manageability interfaces exposed by a manageable Web service and accessible by the management system that assumes the role of the manager. A WSDL document exists that describes the manageability interfaces and the endpoints to which messages with management related payloads could be sent.

Note: An application might assume the role of the client and the role of the manager simultaneously. In this case, the application might exercise business functions of the Web service and manage it at the same time.

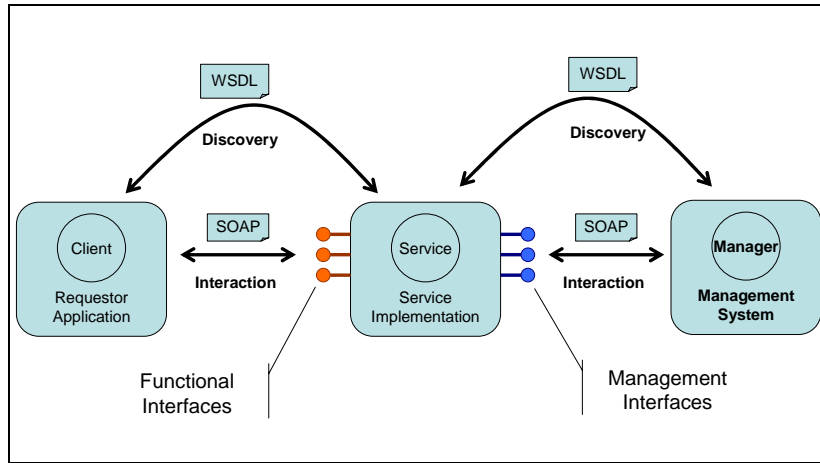


Figure 13: Manageable Web Services

606
607
608
609

4.3. Management Patterns

610
611

4.3.1. Consolidated Interfaces

612
613
614
615
616
617
618
619

In the first management pattern (Figure 14), one WSDL document describes one service with an endpoint whose implementation supports both functional and manageability interfaces. Messages carrying functional requests and management requests can be sent to the same endpoint, and the implementation will do the proper task depending on the content of the message. The endpoint will provide access control and authorization to the operations in the manageability and business interfaces.

620
621
622
623
624
625
626
627
628
629
630
631

This pattern is applicable when a Web service provider intends to offer business functionality and its manageability within a single description of the service. This is the simplest method for implementation and use of the business functionality and its manageability. It is similar to security policies and enforcement mechanisms being described along with the service they apply to. The drawback of this pattern is the “all or nothing” nature it brings to visibility. Whoever discovers and uses the business functionality also discovers the manageability capabilities associated with the service. It does not mean that manageability capabilities can automatically be used by the same clients that use business functionality; access control policies may only allow certain known clients to use the manageability interface.




Figure 14: Consolidated Interfaces Management Pattern

632
633

634

635

636

4.3.2. Associated Interfaces

637

In the second management pattern (Figure 15), the manageability is implemented and offered completely separately from the service itself. For example, it could be a management agent implanted into the hosting environment, or it could be that the environment itself provides manageability capabilities. In this case, two WSDL documents may exist: one that describes the functional Web service endpoint and another one that describes manageability interfaces exposed via an endpoint (manageability endpoint). A referencing mechanism associates the two endpoint descriptions and makes it possible to manage the Web service endpoint. Possible referencing mechanisms are outlined in the [Describing and Discovering Manageability](#) section later in this document.

647

648

649

650

651

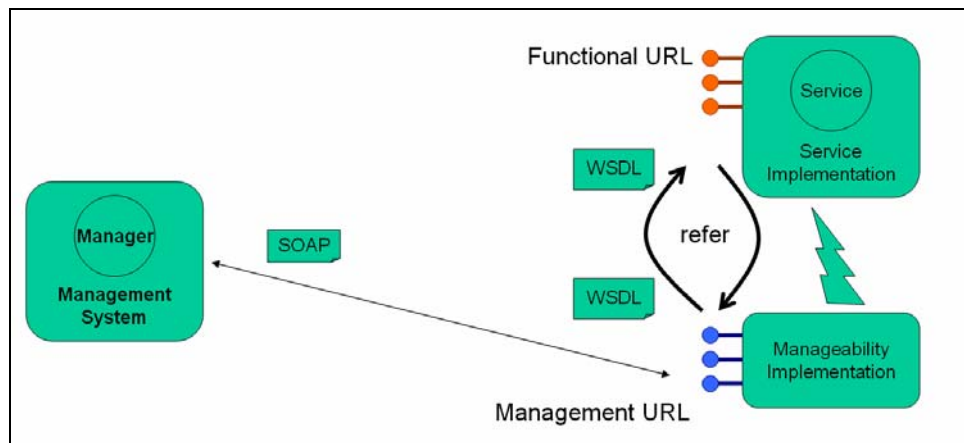
652

653

654

655

This pattern is applicable when a Web service provider needs to separate visibility of the service from its manageability. It may be complex to implement, deploy, and use for the managers, but developers of the service implementation and client applications are not impacted directly with the manageability description and publication. It becomes fully the responsibility of the deployment and management activities. If a client needs to be aware of the manageability, such as it may need to query if service is up, then the client needs to be developed such that it is aware of both interfaces.



656

657

Figure 15: Associated Interfaces Management Pattern

658

659 4.4. Managers and Manageability Interfaces

660

661

4.4.1. Implementing the Manageability Interface

662

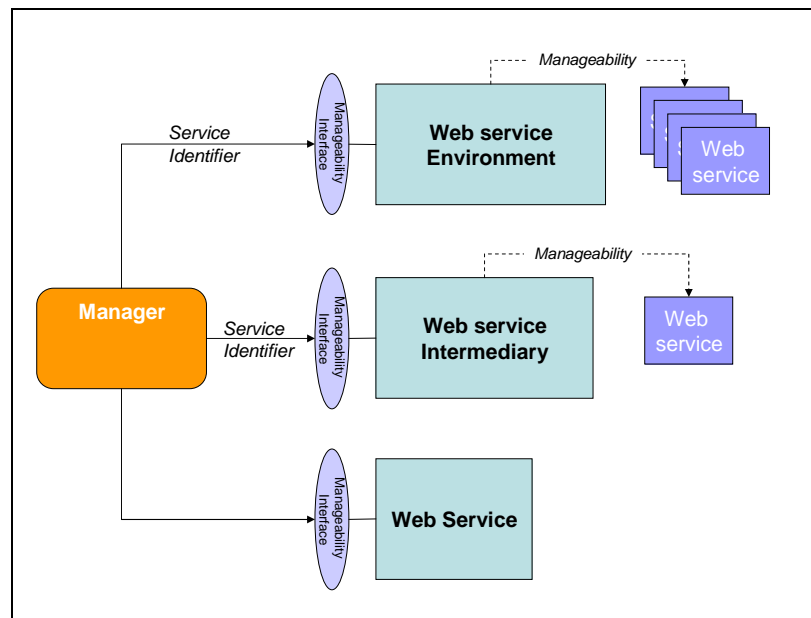
663 It is important that the manageability defined in this document be implementable in
 664 an easy, flexible, and interoperable manner. Many implementation and deployment
 665 techniques must be supported, and manageability must be achievable across many
 666 disparate solutions.

667

668

669 Figure 16 outlines three major methods of making a Web service manageable. First, a
 670 Web service implementation itself can offer manageability interfaces. This would be
 671 an *explicitly manageable* Web service implementation. Second, an *intermediary* can
 672 add the manageability interface to a Web service it represents. An intermediary in this
 673 case is used as a generic term for a service proxy or a service agent that is capable of
 674 interjecting itself between a service and a client. And, finally, an *environment* can
 675 offer manageability interfaces for the Web services it hosts. The latter two methods
 676 *delegate* provision of the manageability of a Web service to that other element (the
 environment or the intermediary).

677



677

678

679

Figure 16: Delegated Manageability within the Web services Architecture

680

4.4.2. Multiple Managers Considerations

681

682 In large-scale enterprises where manageability capabilities and interfaces are
 683 standardized, individual services can be managed by multiple managers. Two
 684 common multiple manager scenarios are (1) separation of responsibilities between
 685 managers (i.e. one does performance monitoring and another one takes care of
 686 deployment changes) and (2) installations with management software from various
 687 vendors. The manageability interface and the manageability model must enable
 operational consistency. The provider of the manageability must ensure consistency

688 of simultaneously executed operations, and managers must ensure there are not
689 executing conflicting operations.

690

691 As outlined earlier ([Scope of Work](#)), this document and the related Web Services
692 Manageability Specification do not prescribe precisely how to implement
693 manageability of a Web service or how to implement a manager. Only the agreed
694 information model, access mechanism, and architecture are described for the purposes
695 of the interoperability between manageable Web services and managers. Although
696 manageability information model itself does not preclude multiple managers, actual
697 implementation of the manageability of a Web service or a manager may or may not
698 be capable of properly supporting multiple simultaneous managers.

699

700 Implementers and product architects supporting multiple manager environments
701 should consider the following recommendations:

702

703 ➤ Proper *access control* to the manageability capabilities of a Web service must be
704 in place. The authentication and policy awareness/enforcement mechanisms must
705 be compliant with Web Services Architecture. For example, a requirement for a
706 WS-Security with X.509 certificates acquired from an identity authority could be
707 stated by the security policy included in the description of the manageability
708 interface (e.g. a WS-Policy statement).

709

710 ➤ An implementation of the manageability capabilities of a Web service and the
711 implementation of the manager must be aware of common and manager specific
712 information in the manageability model. For example, if a Web service endpoint
713 is not accepting any new requests, its state is 'DOWN' for all managers (common
714 data). However, a set of event conditions about transitions or properties which an
715 individual manager is interested in are specific to that manager.

716

717 ➤ An implementation of the manageability capabilities of a Web service should
718 provide a mechanism to establish *affinity* with between the service and the
719 manager. For example, a session/context mechanism such as WS-
720 SecureConversation [7] could be used. A Manager may be required to establish a
721 context that expires over time. Alternatively, a context mechanism described in
722 WS-Coordination [6] may be used or any other preferred standard method.

723

724 ➤ Managers should be capable of making themselves *aware of the changing*
725 *situation* around the manageable Web service. For example, if a particular state
726 transition, configuration property or performance metric is important to a
727 manager, it has to subscribe to events with respect to that information. Since other
728 managers may cause state changes, the manager cannot assume that it will be the
729 origin of all state changes for the manageable Web service.

730

731 ➤ To avoid conflicting operation of the manageability of a Web service, managers
732 must *coordinate and negotiate usage* amongst themselves. Coordination strategies
733 should a) happen among peers – federation based on negotiation, b) happen in a
734 delegated manner – a trusted third party coordinates usage, or c) happen by direct

735 command – a hierarchy of managers. Standards must be used to coordinate and
736 negotiate usage of the manageability. For example, federated coordination may be
737 supported by WS-Agreement [\[5\]](#) and delegated coordination by WS-Coordination
738 [\[6\]](#).
739

740 5. Description and Discovery of Manageability

741 A manageable Web Service (as described in section 4) has to provide machine processable
742 description of the manageability capabilities and the manageability must be discoverable by
743 the managers. Both the description and discovery must be conformant to the Web Services
744 Architecture.

745 5.1. Describing the Manageability

746 The Web Services Management Architecture requires a Web service to have a
747 description of its manageability capabilities in a WSDL document. The WSDL 1.2
748 definition of the principles of describing services, interfaces and endpoints is
749 available at http://www.w3.org/TR/wsdl12/#intro_ws. Similar information about
750 WSDL 1.1 is available http://www.w3.org/TR/wsdl#_introduction.

751
752 The manageability interfaces in this work are described in WSDL 1.1 (or in a WSDL
753 with Grid extensions, i.e. GWSDL (see WS-Manageability Specification). These
754 interfaces have to be made available for (or associated with) a given functional Web
755 Service endpoint. According to the earlier requirement and to the WSDL itself, the
756 manageability interface must also be offered via a Web Service endpoint.

757
758 The principles of describing the manageability of a Web Service and its endpoints are
759 defined following a set of manageability patterns suggested in section 4.3.

760 5.1.1. Consolidated Interfaces

761 WSDL 1.2 draft proposes an interface composition mechanism directly in the
762 language. A WSDL 1.2 fragment that represents consolidation of interfaces may look
763 as follows.
764

```
765 <definitions name="AConsolidationOfManageabilityExampleInWSDL12"
766   targetNamespace="http://www.oasis-open.org/wsdm/example/3"
767   xmlns:tns="http://www.oasis-open.org/wsdm/example/3"
768   xmlns:wsdm="urn:wsdm:webservice:endpoint:manageability" ... >
769
770 <import namespace="urn:wsdm:webservice:endpoint:manageability"
771   location="http://www.oasis-
772   open.org/wsdm/ManageableWebServiceEndpoint.wsdl">
773   [. . .]
774 <interface name="myBusinessFunction"> ... </interface>
775 [. . .]
776 <interface name="myServiceInterface"
777   extends="tns:myBusinessFunction wsdm:ManageableEndpoint"/>
778 ...
779 <service name="myService" interface="tns:myServiceInterface">
780   <endpoint name="myEndpoint" binding="tns:SecureSOAP">
781     <soap:address>http://mycorp.com/myService</soap:address>
782   </endpoint>
783 </service>
784 </definitions>
```

786

787 WSDL 1.1 does not have interface composition mechanism as part of the language.
 788 Therefore, consolidation of interfaces is only possible by composing (appending)
 789 operations of the manageability interface directly into the functional interface. The
 790 appended manageability operations must then be bound to endpoints in the same way
 791 that other functional operations were bound. The following WSDL 1.1 fragment
 792 represents consolidation of interfaces.

```
793
794 <definitions name="AConsolidationOfManageabilityExampleInWSDL11"
795   targetNamespace="http://www.oasis-open.org/wsdm/example/4"
796   xmlns:tns="http://www.oasis-open.org/wsdm/example/4"
797   xmlns:wsdm="urn:wsdm:webservice:endpoint:metrics:data" ... >
798
799 <import namespace=" urn:wsdm:webservice:endpoint:metrics:data"
800   location="http://www.oasis-
801 open.org/wsdm/ManageableWebServiceEndpointMetricsData.wsdl">
802 [. . .]
803 <portType name="myServiceInterface"> ...
804   <operation name="myBusinessOperation"> ... </operation>
805   <operation name="getRequestCounters"> ... </operation> ...
806 </portType>
807 [. . .]
808 <binding name="mySecureSOAP" portType="tns:myServiceInterface"> ...
809   <operation name="myBusinessOperation"> ... </operation>
810   <operation name="getRequestCounters"> ... </operation> ...
811 </binding>
812 [. . .]
813 <service name="myService">
814   <port name="myEndpoint" binding="tns:mySecureSOAP">
815     <soap:address>http://mycorp.com/myService</soap:address>
816   </port>
817 </service>
818 </definitions>
```

819

820 **Note** : In the case of consolidated interfaces there is no need to identify at runtime which endpoint
 821 manageability applies to. It is the same endpoint that exposes both interfaces. The functional
 822 interface and management interface are associated by being defined in the same service
 823 description. They can be defined in the same interface by using inheritance from business and
 824 manageability interfaces or by implementing the operations directly in the endpoint's interface.

825

826 5.1.2. Associated Interfaces

827 There could be various mechanisms to associate the interfaces. Following are the
 828 mechanisms that are most likely to be used.

- 829 ➤ Associating at “design time” by embedding a reference in the WSDL
 830 document.
- 831 ➤ Associating at “runtime” by an interface method that returns a reference.
- 832 ➤ Associating by external “discovery” of manageable elements.

833

834 Associating at runtime and by discovery is outlined in the section 5.2. The method of
 835 association at design time using WSDL extensions is outlined in section 1.2 of the
 836 Web Services Manageability – Relationships document.

837

838 **5.2. Discovery of Manageability**

839 In order for a manager to make use of the manageability of a Web service endpoint, it
840 needs to:

- 841 ➤ Understand that a Web Service endpoint is manageable
- 842 ➤ Understand what manageability capabilities are supported by a Web service
843 endpoint.

844

845 The process of discovery of the manageability is required to be compliant with the
846 Web Services Architecture.

847

848 **5.2.1. Discovery of Manageable Web Services**

849 Manageable Web Service endpoints can be discovered by managers in the following
850 ways.

- 851 ➤ Use any standard Web Services discovery mechanism (e.g. search engine,
852 UDDI registry, peer discovery such as WS-Inspection, etc.). Then locate
853 manageability endpoint (consolidated or associated).
- 854 ➤ Discover manageability endpoint via management discovery mechanism
855 (proprietary or Web Services Architecture compliant, e.g. a UDDI registry).
856 Then request identification of the Web service endpoint it applies to.
- 857 ➤ Follow relationships of a manageable Web service endpoint to locate other
858 related elements. Some of the related elements may be other Web service
859 endpoints which also have manageability endpoints.

860

861 Basically, either way, after discovering a Web service endpoint or manageability
862 endpoint, the manager has to understand which one is which. Following is the brief
863 outline of how to locate the manageability endpoint having discovered a Web service
864 endpoint and vice versa.

865 *From functional Web Service endpoint to its manageability endpoint*

866 When a manager discovers a Web service endpoint it must have the WSDL document
867 that describes it. To determine if the given Web service endpoint is manageable, the
868 manager must interrogate the WSDL document first. It must understand whether
869 consolidated interfaces or associated interfaces patterns are used.

870

871 To determine if consolidated interfaces pattern was used, the manager simply finds if
872 the endpoint offers standard manageability operations. It finds those operations by the
873 input/output message definition QName. For example, a sufficient indication that an
874 endpoint offers standard metrics data is an operation which the input message is
875 **{urn:wsdm:webservice:endpoint:metrics:data,getRequestCountersIn}**. This
876 approach works equally well for WSDL 1.1, GWSDL and 1.2.

877

878 In case of WSDL 1.2 there is a simpler way to determine if an endpoint is
879 manageable in a consolidated manner. The manager could simply look up the
880 “composition” line of an interface which the endpoint offers. The “composition” line
881 is formed by taking wsdl:service element which contains the given endpoint
882 description as wsdl:endpoint element and then recursively looking up wsdl:extends
883 attribute of every interface that the service implements (wsdl:implements on

884 wsdl:service). For example, the manager may be looking for the
885 **{urn:wsdm:webservice:endpoint:manageability,ManageableEndpoint}** interface
886 QName or by inheritance from known manageability interfaces to determine if the
887 endpoint is manageable.

888

889 To determine if associated interfaces pattern was used, a manager could use the
890 WSDL extensibility mechanism described in section WS-Manageability specification
891 document. A few XML tags from **urn:wsrl:webservice:endpoint:relationship**
892 namespace appearing in the WSDL document would point the manager to the
893 manageability endpoint.

894

895 The associated manageability endpoint could be established using the discovery
896 mechanism such as UDDI registry. For example, a registry could contain an entry for
897 a service (business service) and then two implementation entries (tModels). One
898 tModel could contain a reference to a WSDL that describes the functional Web
899 service endpoint, and another could contain a reference to a WSDL that describes the
900 associated manageability Web service endpoint. The registry could control the
901 visibility and only allow appropriate parties to see the second tModel pointing to the
902 manageability endpoint description.

903 *From manageability endpoint to the functional Web Service endpoint*

904 A manager may know about a manageability endpoint. It may have established such
905 an endpoint using a management-oriented discovery or by following relationships of
906 manageable elements, etc. The manager, then, may need to find a description of the
907 Web service endpoint which the manageability is provided for. For example it may
908 need to know what functional operations are available, or it may want to send a
909 synthetic transaction request to the functional endpoint.

910

911 Manageability interface declares access to the identification information of the
912 endpoint it provides manageability for. A manager may use this information to locate
913 the description of the functional Web service endpoint. Access to the identification
914 must always be provided by the manageability endpoint that is used in the associated
915 interfaces pattern.

916

917 Note that in case of consolidated interfaces pattern, manageability endpoint is the
918 same as the functional endpoint and, therefore, may choose not to provide access to
919 the identification information. Determining this is trivial for the manager, though. A
920 manager must assume that if a manageable endpoint does not provide access to the
921 identification information, it is a case of the consolidated interfaces pattern. A
922 manager must then use a standard Web services discovery approach to determine the
923 functional capabilities of the endpoint if that becomes necessary.

924

925 **5.3. Discovery of Manageability Capabilities**

926 A manager may need to know what capabilities a given manageability endpoint
927 offers. There are two ways to determine that: (1) by interrogating the endpoint
928 description, and (2) by asking the endpoint itself. Certain capabilities can be

929 discovered only from the description, and certain others can only be determined by
930 asking (sending request at runtime).

931

932 For example, to determine if a manageability endpoint offers request counters data, a
933 manager would simply lookup the description document if there is an operation with
934 the input message

935 **{urn:wsdm:webservice:endpoint:metrics:data,getRequestCountersIn}** QName.

936 This way most other principal capabilities described in the composable, standard
937 manageability interfaces can be determined.

938

939 An example of when a manager has to ask the manageability endpoint could be when
940 determining the availability of a particular metric. In this case, a manager must send a
941 request to retrieve the list of available metric QNames. It then has to check the
942 returned list if a metric it is interested in is actually provided by the given
943 manageability endpoint.

944

945 **6. Summary**

946 This document has outlined the fundamental concepts of Web services architectures, how
947 the architecture is realized, and how the realization influences the manageability of Web
948 services. Patterns for supporting manageability interfaces and discovery are outlined here
949 to provide guidance during the development of manageability specifications. The WS-
950 Manageability – Specification builds on these concepts to define the manageability model
951 for Web services in an implementation agnostic manner. The WS-Manageability –
952 Representations document renders the manageability model into interfaces expressed in
953 WSDL 1.1 and GWSDL.

954

955

956 7. References

957

958

959

960

961

962

963

964

965

966

967

968

969

- [1] **SOA – Service Oriented Architecture:** http://www.w3.org/TR/2003/WD-ws-arch-20030808/#service_oriented_architecture
- [2] **Grid:** <http://www.globus.org/research/papers.html#Overview%20Papers> -
- [3] **Web Services Architecture:** <http://www.w3.org/TR/wsa-reqs>
- [4] **WS-I Basic Profile:** <http://www.ws-i.org/Profiles/Basic/2003-08/BasicProfile-1.0a.htm> :
- [5] **WS-Agreement:**
http://www.globus.org/research/papers/OGSI_Agreement_2003_06_12.pdf
- [6] **WS-Coordination:** <http://www-106.ibm.com/developerworks/library/ws-coor/> :
- [7] **WS-SecureConversation:** <http://www-106.ibm.com/developerworks/library/ws-secon/>