



Web Services Atomic Transaction (WS-AtomicTransaction) 1.1

Public Review Draft 01, August 30, 2006

Document Identifier:

wstx-wsat-1.1-spec-pr-01

Location:

<http://docs.oasis-open.org/ws-tx/wstx-wsat-1.1-spec-pr-01.pdf>

Technical Committee:

OASIS WS-TX TC

Chair(s):

Eric Newcomer, Iona
Ian Robinson, IBM

Editor(s):

Mark Little, JBoss Inc. <mark.little@jboss.com>
Andrew Wilkinson, IBM <awilkinson@uk.ibm.com>

Abstract:

This specification provides the definition of the atomic transaction coordination type that is to be used with the extensible coordination framework described in the WS-Coordination specification. The specification defines three specific agreement coordination protocols for the atomic transaction coordination type: completion, volatile two-phase commit, and durable two-phase commit. Developers can use any or all of these protocols when building applications that require consistent agreement on the outcome of short-lived distributed activities that have the all-or-nothing property.

Status:

This document is published by the WS-TX TC as a "public review draft".

This document was last revised or approved by the WS-TX TC on the above date. The level of approval is also listed above. Check the current location noted above for possible later revisions of this document. This document is updated periodically on no particular schedule.

Technical Committee members should send comments on this specification to the Technical Committee's email list. Others should send comments to the Technical Committee by using the "Send A Comment" button on the Technical Committee's web page at www.oasis-open.org/committees/ws-tx.

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

The non-normative errata page for this specification is located at www.oasis-open.org/committees/ws-tx.

Notices

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

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

Copyright © OASIS Open 2006. *All Rights Reserved.*

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

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

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

Table of contents

1	Introduction	4
1.1	Composable Architecture.....	4
1.2	Terminology	4
1.3	Namespace	5
1.3.1	Prefix Namespace	5
1.4	XSD and WSDL Files.....	5
1.5	AT Protocol Elements	6
1.6	Normative References	6
2	Atomic Transaction Context.....	8
3	Atomic Transaction Protocols.....	9
3.1	Preconditions	9
3.2	Completion Protocol.....	9
3.3	Two-Phase Commit Protocol	10
3.3.1	Volatile Two-Phase Commit Protocol.....	10
3.3.2	Durable Two-Phase Commit Protocol	11
3.3.3	2PC Diagram and Notifications.....	11
4	AT Policy Assertion.....	13
4.1	Assertion Model	13
4.2	Normative Outline	13
4.3	Assertion Attachment.....	13
4.4	Assertion Example	13
5	Transaction Faults	15
5.1	Inconsistent Internal State.....	16
5.2	Unknown Transaction	16
6	Security Model	17
7	Security Considerations.....	19
8	Use of WS-Addressing Headers.....	21
9	State Tables.....	22
9.1	Completion Protocol.....	22
9.2	2PC Protocol	23
A.	Acknowledgements.....	25
B.	Revision History	26

1 Introduction

The current set of Web service specifications [[WSDL](#)][[SOAP11](#)][[SOAP12](#)] defines protocols for Web service interoperability. Web services increasingly tie together a number of participants forming large distributed applications. The resulting activities may have complex structure and relationships.

The WS-Coordination [[WSCOOR](#)] specification defines an extensible framework for defining coordination types. This specification provides the definition of an atomic transaction coordination type used to coordinate activities having an "all or nothing" property. Atomic transactions commonly require a high level of trust between participants and are short in duration. The Atomic Transaction specification defines protocols that enable existing transaction processing systems to wrap their proprietary protocols and interoperate across different hardware and software vendors.

To understand the protocol described in this specification, the following assumptions are made:

- The reader is familiar with existing standards for two-phase commit protocols and with commercially available implementations of such protocols. Therefore this section includes only those details that are essential to understanding the protocols described.
- The reader is familiar with the WS-Coordination specification that defines the framework for the WS-AtomicTransaction coordination protocols.
- The reader is familiar with WS-Addressing [[WSADDR](#)] and WS-Policy [[WSPOLICY](#)].

Atomic transactions have an all-or-nothing property. The actions taken prior to commit are only tentative (i.e., not persistent and not visible to other activities). When an application finishes, it requests the coordinator to determine the outcome for the transaction. The coordinator determines if there were any processing failures by asking the participants to vote. If the participants all vote that they were able to execute successfully, the coordinator commits all actions taken. If a participant votes that it needs to abort or a participant does not respond at all, the coordinator aborts all actions taken. Commit makes the tentative actions visible to other transactions. Abort makes the tentative actions appear as if the actions never happened. Atomic transactions have proven to be extremely valuable for many applications. They provide consistent failure and recovery semantics, so the applications no longer need to deal with the mechanics of determining a mutually agreed outcome decision or to figure out how to recover from a large number of possible inconsistent states.

Atomic Transaction defines protocols that govern the outcome of atomic transactions. It is expected that existing transaction processing systems wrap their proprietary mechanisms and interoperate across different vendor implementations.

1.1 Composable Architecture

By using the XML [[XML](#)], SOAP [[SOAP11](#)][[SOAP12](#)] and WSDL [[WSDL](#)] extensibility model, SOAP-based and WSDL-based specifications are designed to work together to define a rich Web services environment. As such, WS-AtomicTransaction by itself does not define all features required for a complete solution. WS-AtomicTransaction is a building block used with other specifications of Web services (e.g., WS-Coordination, WS-Security) and application-specific protocols that are able to accommodate a wide variety of coordination protocols related to the coordination actions of distributed applications.

1.2 Terminology

The uppercase key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC2119 [[KEYWORDS](#)].

Namespace URIs of the general form "some-URI" represents some application-dependent or context-dependent URI as defined in RFC3986 [[URI](#)].

46 This specification uses an informal syntax to describe the XML grammar of the XML fragments below:

47 • The syntax appears as an XML instance, but the values indicate the data types instead of values.

48 • Element names ending in "..." (such as <element.../> or <element...>) indicate that

49 elements/attributes irrelevant to the context are being omitted.

50 • Attributed names ending in "..." (such as name=...) indicate that the values are specified below.

51 • Grammar in bold has not been introduced earlier in the document, or is of particular interest in an

52 example.

53 • <!-- description --> is a placeholder for elements from some "other" namespace (like ##other in XSD).

54 • Characters are appended to elements, attributes, and <!-- descriptions --> as follows: "?" (0 or 1), "*" (0 or more), "+" (1 or more). The characters "[" and "]" are used to indicate that contained items are to

55 be treated as a group with respect to the "?", "*", or "+" characters.

56

57 • The XML namespace prefixes (defined below) are used to indicate the namespace of the element

58 being defined.

59 • Examples starting with <?xml contain enough information to conform to this specification; others

60 examples are fragments and require additional information to be specified in order to conform.

61 XSD schemas and WSDL definitions are provided as a formal definition of grammars [[XML-Schema1](#)]

62 [[WSDL](#)].

63 1.3 Namespace

64 The XML namespace URI that MUST be used by implementations of this specification is:

```
65 http://docs.oasis-open.org/ws-tx/wsat/2006/06
```

66 This is also used as the CoordinationContext type for atomic transactions.

67 1.3.1 Prefix Namespace

Prefix	Namespace
S11	http://schemas.xmlsoap.org/soap/envelope
S12	http://www.w3.org/2003/05/soap-envelope
wscor	http://docs.oasis-open.org/ws-tx/wscor/2006/06
wsat	http://docs.oasis-open.org/ws-tx/wsathat/2006/06

68 If an action URI is used then the action URI MUST consist of the wsat namespace URI concatenated with

69 the "/" character and the element name. For example:

```
70 http://docs.oasis-open.org/ws-tx/wsathat/2006/06/Commit
```

71 1.4 XSD and WSDL Files

72 The following links hold the XML schema and the WSDL declarations defined in this document.

73 <http://docs.oasis-open.org/ws-tx/wsathat/2006/06/wsathat.xsd>

74 <http://docs.oasis-open.org/ws-tx/wsathat/2006/06/wsathat.wsdl>

75 SOAP bindings for the WSDL documents defined in this specification MUST use "document" for the *style*

76 attribute.

77 **1.5 AT Protocol Elements**

78 The protocol elements define various extensibility points that allow other child or attribute content.
79 Additional children and/or attributes MAY be added at the indicated extension points but MUST NOT
80 contradict the semantics of the parent and/or owner, respectively. If a receiver does not recognize an
81 extension, the receiver SHOULD ignore the extension.

82 **1.6 Normative References**

83 **[KEYWORDS]**

84 S. Bradner, "Key words for use in RFCs to Indicate Requirement Levels", RFC 2119,
85 <http://www.ietf.org/rfc/rfc2119.txt>, Harvard University, March 1997

86 **[SOAP11]**

87 W3C Note, "SOAP: Simple Object Access Protocol 1.1", [http://www.w3.org/TR/2000/NOTE-](http://www.w3.org/TR/2000/NOTE-SOAP-20000508)
88 [SOAP-20000508](http://www.w3.org/TR/2000/NOTE-SOAP-20000508), 08 May 2000

89 **[SOAP12]**

90 W3C Recommendation, "SOAP Version 1.2 Part 1: Messaging Framework",
91 <http://www.w3.org/2003/05/soap-envelope>, June 2003

92 **[URI]**

93 T. Berners-Lee, R. Fielding, L. Masinter, "Uniform Resource Identifiers (URI): Generic Syntax",
94 RFC 3986, <http://www.ietf.org/rfc/rfc3986.txt>, MIT/LCS, Day Software, Adobe Systems, January
95 2005

96 **[WSADDR]**

97 Web Services Addressing (WS-Addressing) 1.0, <http://www.w3.org/2005/08/addressing>, W3C
98 Recommendation, May 2006

99 **[WSCOOR]**

100 Web Services Coordination (WS-Coordination) 1.1, [http://docs.oasis-open.org/ws-](http://docs.oasis-open.org/ws-tx/wscoor/2006/06)
101 [tx/wscoor/2006/06](http://docs.oasis-open.org/ws-tx/wscoor/2006/06), OASIS, March 2006

102 **[WSDL]**

103 Web Services Description Language (WSDL) 1.1, [http://www.w3.org/TR/2001/NOTE-wsdl-](http://www.w3.org/TR/2001/NOTE-wsdl-20010315)
104 [20010315](http://www.w3.org/TR/2001/NOTE-wsdl-20010315)

105 **[WSPOLICY]**

106 Web Services Policy Framework (WS-Policy), <http://schemas.xmlsoap.org/ws/2004/09/policy>,
107 VeriSign, Microsoft, Sonic Software, IBM, BEA Systems, SAP, September 2004

108 **[WSPOLICYATTACH]**

109 Web Services Policy Attachment (WS-PolicyAttachment),
110 <http://schemas.xmlsoap.org/ws/2004/09/policy>, VeriSign, Microsoft, Sonic Software, IBM, BEA
111 Systems, SAP, September 2004

112 **[WSSec]**

113 OASIS Standard 200401, "Web Services Security: SOAP Message Security 1.0 (WS-Security
114 2004)", [http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-soap-message-security-](http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-soap-message-security-1.0.pdf)
115 [1.0.pdf](http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-soap-message-security-1.0.pdf), March 2004

116 **[WSSecConv]**

117 Web Services Secure Conversation Language (WS-SecureConversation),
118 <http://schemas.xmlsoap.org/ws/2005/02/sc>, OpenNetwork, Layer7, Netegrity, Microsoft,

- 119 Reactivity, IBM, VeriSign, BEA Systems, Oblix, RSA Security, Ping Identity, Westbridge,
120 Computer Associates, February 2005
- 121 **[WSecPolicy]**
- 122 Web Services Security Policy Language (WS-SecurityPolicy),
123 <http://schemas.xmlsoap.org/ws/2005/07/securitypolicy>, Microsoft, VeriSign, IBM, RSA Security,
124 July 2005
- 125 **[WSTrust]**
- 126 Web Services Trust Language (WS-Trust), , <http://schemas.xmlsoap.org/ws/2005/02/trust>,
127 OpenNetwork, Layer7, Netegrity, Microsoft, Reactivity, VeriSign, IBM, BEA Systems, Oblix, RSA
128 Security, Ping Identity, Westbridge, Computer Associates, February 2005
- 129 **[XML]**
- 130 W3C Recommendation, "Extensible Markup Language (XML) 1.0 (Fourth Edition)",
131 <http://www.w3.org/TR/2006/REC-xml-20060816>, 16 August 2006
- 132 **[XML-ns]**
- 133 W3C Recommendation, "Namespaces in XML (Second Edition)",
134 <http://www.w3.org/TR/2006/REC-xml-names-20060816>, 16 August 2006
- 135 **[XML-Schema1]**
- 136 W3C Recommendation, " XML Schema Part 1: Structures Second Edition",
137 <http://www.w3.org/TR/2004/REC-xmlschema-1-20041028>, 28 October 2004
- 138 **[XML-Schema2]**
- 139 W3C Recommendation, " XML Schema Part 2: Datatypes Second Edition",
140 <http://www.w3.org/TR/2004/REC-xmlschema-2-20041028>, 28 October 2004

141 2 Atomic Transaction Context

142 Atomic Transaction builds on WS-Coordination, which defines an activation and a registration service.
143 Example message flows and a complete description of creating and registering for coordinated activities
144 is found in the WS-Coordination specification [WSCOOR].

145 | The Atomic Transaction coordination context **MUST** flow on all application messages involved with the
146 transaction.

Deleted: must

147 Atomic Transaction adds the following semantics to the CreateCoordinationContext operation on the
148 activation service.

- 149 • If the request includes the CurrentContext element, the target coordinator is interposed as a
150 subordinate to the coordinator stipulated inside the CurrentContext element.
- 151 • If the request does not include a CurrentContext element, the target coordinator creates a new
152 transaction and acts as the root.

153 A coordination context MAY have an Expires element. This element specifies the period, measured from
154 the point in time at which the context was first created or received, after which a transaction MAY be
155 terminated solely due to its length of operation. From that point forward, the coordinator MAY elect to
156 unilaterally roll back the transaction, so long as it has not made a commit decision. Similarly a 2PC
157 participant MAY elect to abort its work in the transaction so long as it has not already decided to prepare.

158 The Atomic Transaction protocol is identified by the following coordination type:

159 `http://docs.oasis-open.org/ws-tx/wsat/2006/06`

3 Atomic Transaction Protocols

160
161
162
163
164
165
166
167
168
169
170
171
172

173
174
175
176
177
178
179

180
181
182
183
184
185
186
187
188
189

This specification defines the following protocols for atomic transactions.

- **Completion:** The completion protocol initiates commitment processing. Based on each protocol's registered participants, the coordinator begins with Volatile 2PC then proceeds through Durable 2PC. The final result is signaled to the initiator.
- **Two-Phase Commit (2PC):** The 2PC protocol coordinates registered participants to reach a commit or abort decision, and ensures that all participants are informed of the final result. The 2PC protocol has two variants:
 - **Volatile 2PC:** Participants managing volatile resources such as a cache should register for this protocol.
 - **Durable 2PC:** Participants managing durable resources such as a database should register for this protocol.

A participant can register for more than one of these protocols by sending multiple Register messages.

3.1 Preconditions

The correct operation of the protocols requires that a number of preconditions **MUST** be established prior to the processing:

1. The source **SHOULD** have knowledge of the destination's policies, if any, and the source **SHOULD** be capable of formulating messages that adhere to this policy.
2. If a secure exchange of messages is required, then the source and destination **MUST** have a security context.

Deleted: MUST
Deleted: MUST

3.2 Completion Protocol


The Completion protocol is used by an application to tell the coordinator to either try to commit or abort an atomic transaction. After the transaction has completed, a status is returned to the application.

An initiator registers for this protocol using the following protocol identifier:

```
http://docs.oasis-open.org/ws-tx/wsac/2006/06/Completion
```

A Completion protocol coordinator must be the root coordinator of an atomic transaction. The registration service for a subordinate coordinator **MUST** respond to an attempt to register for this coordination protocol with the WS-Coordination fault Cannot Register Participant.

The diagram below illustrates the protocol abstractly. Refer to the section State Tables for a detailed description of this protocol.



190

191 The coordinator accepts:

192 Commit

193 Upon receipt of this notification, the coordinator knows that the participant has completed
 194 application processing and that it **SHOULD** attempt to commit the transaction.

Deleted: should

195 Rollback

196 Upon receipt of this notification, the coordinator knows that the participant has terminated
 197 application processing and that it **MUST** abort the transaction.

Deleted: should

198 The initiator accepts:

199 Committed

200 Upon receipt of this notification, the initiator knows that the coordinator reached a decision to
 201 commit.

202 Aborted

203 Upon receipt of this notification, the initiator knows that the coordinator reached a decision to
 204 abort.

205 A coordination service that supports an Activation service **MUST** support the Completion protocol.

206 3.3 Two-Phase Commit Protocol

207 The Two-Phase Commit (2PC) protocol is a Coordination protocol that defines how multiple participants
 208 reach agreement on the outcome of an atomic transaction. The 2PC protocol has two variants: Durable
 209 2PC and Volatile 2PC.

210 3.3.1 Volatile Two-Phase Commit Protocol

211 Upon receiving a Commit notification in the completion protocol, the root coordinator begins the prepare
 212 phase of all participants registered for the Volatile 2PC protocol. All participants registered for this
 213 protocol must respond before a Prepare is issued to a participant registered for Durable 2PC. Further
 214 participants may register with the coordinator until the coordinator issues a Prepare to any durable
 215 participant. Once this has happened the Registration Service for the coordinator **MUST** respond to any
 216 further Register requests with a Cannot Register Participant fault message. A volatile recipient is not
 217 guaranteed to receive a notification of the transaction's outcome.

218 Participants register for this protocol using the following protocol identifier:

219 <http://docs.oasis-open.org/ws-tx/wsac/2006/06/Volatile2PC>

220 **3.3.2 Durable Two-Phase Commit Protocol**

221 After receiving a Commit notification in the completion protocol and upon successfully completing the
 222 prepare phase for Volatile 2PC participants, the root coordinator begins the Prepare phase for Durable
 223 2PC participants. All participants registered for this protocol **MUST** respond Prepared or ReadOnly
 224 before a Commit notification is issued to a participant registered for either protocol.


Deleted: must

225 Participants register for this protocol using the following protocol identifier:

```
226 http://docs.oasis-open.org/ws-tx/wsac/2006/06/Durable2PC
```

227 **3.3.3 2PC Diagram and Notifications**

228 The diagram below illustrates the protocol abstractly. Refer to the section State Tables for a detailed
 229 description of this protocol.



230
 231 The participant accepts:

232 Prepare

233 Upon receipt of this notification, the participant knows to enter phase 1 and vote on the outcome
 234 of the transaction. If the participant does not know of the transaction, it **MUST** vote to abort. If
 235 the participant has already voted, it **MUST** resend the same vote.

Deleted: must

Deleted: should

236 Rollback

237 Upon receipt of this notification, the participant knows to abort, and forget, the transaction. This
 238 notification **MAY** be sent in either phase 1 or phase 2. Once sent, the coordinator **MAY** forget all
 239 knowledge of this transaction.

Deleted: can

Deleted: may

240 Commit

241 Upon receipt of this notification, the participant knows to commit the transaction. This notification
 242 **MUST** only be sent after phase 1 and if the participant voted to commit. If the participant does
 243 not know of the transaction, it **MUST** send a Committed notification to the coordinator.

Deleted: can

Deleted: must

244 The coordinator accepts:

245 Prepared

246 Upon receipt of this notification, the coordinator knows the participant is prepared and votes to
 247 commit the transaction.

248 ReadOnly

249 Upon receipt of this notification, the coordinator knows the participant votes to commit the
250 transaction, and has forgotten the transaction. The participant does not wish to participate in
251 phase 2.

252 Aborted

253 Upon receipt of this notification, the coordinator knows the participant has aborted, and forgotten,
254 the transaction.

255 Committed

256 Upon receipt of this notification, the coordinator knows the participant has committed the
257 transaction. That participant ~~MAY~~ be safely forgotten.

Deleted: may

258 Conforming implementations MUST implement the 2PC protocol.

259 4 AT Policy Assertion

260 WS-Policy Framework [WSPOLICY] and WS-Policy Attachment [WSPOLICYATTACH] collectively define
261 a framework, model and grammar for expressing the capabilities, requirements, and general
262 characteristics of entities in an XML Web services-based system. To enable a web service to describe
263 transactional capabilities and requirements of a service and its operations, this specification defines a pair
264 of Atomic Transaction policy assertions that leverage the WS-Policy framework.

265 4.1 Assertion Model

266 The AT policy assertion is provided by a web service to qualify the transactional processing of messages
267 associated with the particular operation to which the assertion is scoped. The AT policy assertion
268 indicates whether a requester MAY or MUST include an AtomicTransaction CoordinationContext flowed
269 with the message.

270 4.2 Normative Outline

271 The normative outline for the AT policy assertion is:

```
272 <wsat:ATAssertion [wsp:Optional="true"]? ... >  
273 ...  
274 </wsat:ATAssertion>
```

275 The following describes additional, normative constraints on the outline listed above:

276 /wsat:ATAssertion

277 A policy assertion that specifies that an atomic transaction MUST be flowed inside a requester's
278 message. From the perspective of the requester, the target service that processes the transaction MUST
279 behave as if it had participated in the transaction. The transaction MUST be represented as a SOAP
280 header in CoordinationContext format, as defined in WS-Coordination [WSCOOR].

281 /wsat:ATAssertion/@wsp:Optional="true"

282 Per WS-Policy [WSPOLICY], this is compact notation for two policy alternatives, one with and one without
283 the assertion.

284 4.3 Assertion Attachment

285 Because the AT policy assertion indicates atomic transaction behavior for a single operation, the
286 assertion has Operation Policy Subject [WSPOLICYATTACH].

287 WS-PolicyAttachment defines two WSDL [WSDL] policy attachment points with Operation Policy Subject:

- 288 • wsdl:portType/wsdl:operation – A policy expression containing the AT policy assertion MUST
289 NOT be attached to a wsdl:portType; the AT policy assertion specifies a concrete behavior
290 whereas the wsdl:portType is an abstract construct.
- 291 • wsdl:binding/wsdl:operation – A policy expression containing the AT policy assertion SHOULD be
292 attached to a wsdl:binding.

293 4.4 Assertion Example

294 An example use of the AT policy assertion follows:

```
295 (01) <wsdl:definitions  
296 (02)     targetNamespace="bank.example.com"
```

```
297 (03)      xmlns:tns="bank.example.com"
298 (04)      xmlns:wSDL="http://schemas.xmlsoap.org/wSDL/"
299 (05)      xmlns:wsp="http://schemas.xmlsoap.org/ws/2004/09/policy"
300 (06)      xmlns:wsat="http://docs.oasis-open.org/ws-tx/wsat/2006/06"
301 (07)      xmlns:wsu="http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-
302 wssecurity-utility-1.0.xsd" >
303 (08)      <wsp:Policy wsu:Id="TransactedPolicy" >
304 (09)          <wsat:ATAssertion wsp:optional="true" />
305 (10)          <!-- omitted assertions -->
306 (11)      </wsp:Policy>
307 (12)      <!-- omitted elements -->
308 (13)      <wSDL:binding name="BankBinding" type="tns:BankPortType" >
309 (14)          <!-- omitted elements -->
310 (15)          <wSDL:operation name="TransferFunds" >
311 (16)              <wsp:PolicyReference URI="#TransactedPolicy" wSDL:required="true"
312 />
313 (17)              <!-- omitted elements -->
314 (18)          </wSDL:operation>
315 (19)      </wSDL:binding>
316 (20) </wSDL:definitions>
```

317

318 Lines (8-11) are a policy expression that includes an AT policy assertion (Line 10) to indicate that an
319 atomic transaction in WS-Coordination [[WSCOOR](#)] format MAY be used.

320 Lines (13-19) are a WSDL [[WSDL](#)] binding. Line (17) indicates that the policy in Lines (9-12) applies to
321 this binding, specifically indicating that an atomic transaction MAY flow inside messages.

5 Transaction Faults

322

323 WS-AtomicTransaction faults MUST include as the [action] property the following fault action URI:

324

```
http://docs.oasis-open.org/ws-tx/wsat/2006/06/fault
```

325 The protocol faults defined in this section are generated if the condition stated in the preamble is met.
326 These faults are targeted at a destination endpoint according to the protocol fault handling rules defined
327 for that protocol.

328 The definitions of faults in this section use the following properties:

329 [Code] The fault code.

330 [Subcode] The fault subcode.

331 [Reason] The English language reason element.

332 [Detail] The detail element. If absent, no detail element is defined for the fault.

333 For SOAP 1.2, the [Code] property MUST be either "Sender" or "Receiver". These properties are
334 serialized into text XML as follows:

335

SOAP Version	Sender	Receiver
SOAP 1.2	S12:Sender	S12:Receiver

336

337 The properties above bind to a SOAP 1.2 fault as follows:

```
338 <S12:Envelope>  
339 <S12:Header>  
340 <wsa:Action>  
341 http://docs.oasis-open.org/ws-tx/wsat/2006/06/fault  
342 </wsa:Action>  
343 <!-- Headers elided for clarity. -->  
344 </S12:Header>  
345 <S12:Body>  
346 <S12:Fault>  
347 <S12:Code>  
348 <S12:Value>[Code]</S12:Value>  
349 <S12:Subcode>  
350 <S12:Value>[Subcode]</S12:Value>  
351 </S12:Subcode>  
352 </S12:Code>  
353 <S12:Reason>  
354 <S12:Text xml:lang="en">[Reason]</S12:Text>  
355 </S12:Reason>  
356 <S12:Detail>  
357 [Detail]  
358 ...  
359 </S12:Detail>  
360 </S12:Fault>  
361 </S12:Body>  
362 </S12:Envelope>
```

363 The properties bind to a SOAP 1.1 fault as follows:

```
364 <S11:Envelope>  
365 <S11:Body>  
366 <S11:Fault>
```

```
367 <faultcode>[Subcode]</faultcode>  
368 <faultstring xml:lang="en">[Reason]</faultstring>  
369 </S11:Fault>  
370 </S11:Body>  
371 </S11:Envelope>
```

372 5.1 Inconsistent Internal State

373 This fault is sent by a participant or coordinator to indicate that a protocol violation has been detected
374 after it is no longer possible to change the outcome of the transaction. This is indicative of a global
375 consistency failure and is an unrecoverable condition.

376 Properties:

377 **[Code]** Sender

378 **[Subcode]** wsat:InconsistentInternalState

379 **[Reason]** A global consistency failure has occurred. This is an unrecoverable condition.

380 **[Detail]** unspecified

381 5.2 Unknown Transaction

382 This fault is sent by a coordinator to indicate that it has no knowledge of the transaction and consequently
383 cannot convey the outcome.

384 Properties:

385 **[Code]** Sender

386 **[Subcode]** wsat:UnknownTransaction

387 **[Reason]** The coordinator has no knowledge of the transaction. This is an unrecoverable condition.

388 **[Detail]** unspecified

389 **6 Security Model**


390 The security model for atomic transactions builds on the model defined in WS-Coordination [WSCOOR].
391 That is, services have policies specifying their requirements and requestors provide claims (either implicit
392 or explicit) and the requisite proof of those claims. Coordination context creation establishes a base
393 secret which can be delegated by the creator as appropriate.

394 Because atomic transactions represent a specific use case rather than the general nature of coordination
395 contexts, additional aspects of the security model can be specified.

396 All access to atomic transaction protocol instances is on the basis of identity. The nature of transactions,
397 specifically the uncertainty of systems means that the security context established to register for the
398 protocol instance may not be available for the entire duration of the protocol.


399 Consider for example the scenarios where a participant has committed its part of the transaction, but for
400 some reason the coordinator never receives acknowledgement of the commit. The result is that when
401 communication is re-established in the future, the coordinator will attempt to confirm the commit status of
402 the participant, but the participant, having committed the transaction and forgotten all information
403 associated with it, no longer has access to the special keys associated with the token.

404 The participant can only prove its identity to the coordinator when it indicates that the specified
405 transaction is not in its log and assumed committed. This is illustrated in the figure below:



406
407 There are, of course, techniques to mitigate this situation but such options will not always be successful.
408 Consequently, when dealing with atomic transactions, it is critical that identity claims always be proven to
409 ensure that correct access control is maintained by coordinators.

410 There is still value in coordination context-specific tokens because they offer a bootstrap mechanism so
411 that all participants need not be pre-authorized. As well, it provides additional security because only
412 those instances of an identity with access to the token will be able to securely interact with the coordinator
413 (limiting privileges strategy). This is illustrated in the figure below:



414
 415 The "list" of authorized participants ensures that application messages having a coordination context are
 416 properly authorized since altering the coordination context ID will not provide additional access unless (1)
 417 the bootstrap key is provided, or (2) the requestor is on the authorized participant "list" of identities.

418

7 Security Considerations

419

It is strongly RECOMMENDED that the communication between services be secured using the mechanisms described in WS-Security [WSSec]. In order to properly secure messages, the body and all relevant headers need to be included in the signature. Specifically, the

421

<wscoor:CoordinationContext> header needs to be signed with the body and other key message

422

headers in order to "bind" the two together.

424

In the event that a participant communicates frequently with a coordinator, it is RECOMMENDED that a security context be established using the mechanisms described in WS-Trust [WSTrust] and WS-SecureConversation [WSSecConv] allowing for potentially more efficient means of authentication.

427

It is common for communication with coordinators to exchange multiple messages. As a result, the usage profile is such that it is susceptible to key attacks. For this reason it is strongly RECOMMENDED that the keys be changed frequently. This "re-keying" can be effected a number of ways. The following list outlines four common techniques:

428

429

430

431

- Attaching a nonce to each message and using it in a derived key function with the shared secret

432

- Using a derived key sequence and switch "generations"

433

- Closing and re-establishing a security context (not possible for delegated keys)

434

- Exchanging new secrets between the parties (not possible for delegated keys)

435

It should be noted that the mechanisms listed above are independent of the SCT and secret returned

436

when the coordination context is created. That is, the keys used to secure the channel may be

437

independent of the key used to prove the right to register with the activity.

438

The security context MAY be re-established using the mechanisms described in WS-Trust [WSTrust] and

439

WS-SecureConversation [WSSecConv]. Similarly, secrets can be exchanged using the mechanisms

440

described in WS-Trust. Note, however, that the current shared secret SHOULD NOT be used to encrypt

441

the new shared secret. Derived keys, the preferred solution from this list, can be specified using the

442

mechanisms described in WS-SecureConversation.

443

The following list summarizes common classes of attacks that apply to this protocol and identifies the mechanism to prevent/mitigate the attacks:

444

445

- **Message alteration** – Alteration is prevented by including signatures of the message information using WS-Security [WSSec].

446

447

- **Message disclosure** – Confidentiality is preserved by encrypting sensitive data using WS-Security.

448

449

- **Key integrity** – Key integrity is maintained by using the strongest algorithms possible (by comparing secured policies – see WS-Policy [WSPOLICY] and WS-SecurityPolicy

450

[WSSecPolicy]).

451

452

- **Authentication** – Authentication is established using the mechanisms described in WS-Security and WS-Trust [WSTrust]. Each message is authenticated using the mechanisms described in

453

WS-Security [WSSec].

454

455

- **Accountability** – Accountability is a function of the type of and string of the key and algorithms

456

being used. In many cases, a strong symmetric key provides sufficient accountability. However,

457

in some environments, strong PKI signatures are required.

458

459

- **Availability** – Many services are subject to a variety of availability attacks. Replay is a common

460

attack and it is RECOMMENDED that this be addressed as described in the next bullet. Other

461

attacks, such as network-level denial of service attacks are harder to avoid and are outside the

462

scope of this specification. That said, care should be taken to ensure that minimal processing be performed prior to any authenticating sequences.

- 463
- 464
- 465
- 466
- **Replay** – Messages may be replayed for a variety of reasons. To detect and eliminate this attack, mechanisms should be used to identify replayed messages such as the timestamp/nonce outlined in WS-Security [[WSSec](#)]. Alternatively, and optionally, other technologies, such as sequencing, can also be used to prevent replay of application messages.

467 8 Use of WS-Addressing Headers

468 The protocols defined in WS-AtomicTransaction use a "one way" message exchange pattern consisting of
469 a sequence of notification messages between a Coordinator and a Participant. There are two types of
470 notification messages used in these protocols:

- 471 • A notification message is a terminal message when it indicates the end of a
472 coordinator/participant relationship. **Committed**, **Aborted** and **ReadOnly** are terminal
473 messages, as are the protocol faults defined in this specification and in [WSCOOR].
- 474 • A notification message is a non-terminal message when it does not indicate the end of a
475 coordinator/participant relationship. **Commit**, **Rollback**, **Prepare** and **Prepared** are non-
476 terminal messages.

477 The following statements define addressing interoperability requirements for the WS-AtomicTransaction
478 message types:

479 Non-terminal notification messages

- 480 • MUST include a [source endpoint] property whose [address] property is not set to
481 'http://www.w3.org/2005/08/addressing/anonymous' or 'http://www.w3.org/2005/08/addressing/none'.

482 Both terminal and non-terminal notification messages

- 483 • MUST include a [reply endpoint] property whose [address] property is set to
484 'http://www.w3.org/2005/08/addressing/none'.

485 Notification messages are addressed by both coordinators and participants using the Endpoint
486 References initially obtained during the Register-RegisterResponse exchange. If a [source endpoint]
487 property is present in a notification message, it MAY be used by the recipient. For example, in cases
488 where a Coordinator or Participant has forgotten a transaction that is completed and needs to respond to
489 a resent protocol message, the [source endpoint] property ~~SHOULD~~ be used as described in section 3.3
490 of WS-Addressing 1.0 – Core [WSADDR]. Permanent loss of connectivity between a coordinator and a
491 participant in an in-doubt state can result in data corruption.

Deleted: should

492 Protocol faults raised by a Coordinator or Participant during the processing of a notification message are
493 terminal notifications and MUST be composed using the same mechanisms as other terminal notification
494 messages.

495 All messages are delivered using connections initiated by the sender.

496 **9 State Tables**

497 The following state tables specify the behavior of coordinators and participants when presented with
 498 protocol messages or internal events.

499 Each cell in the tables uses the following convention:

500

Legend
<i>Action to take</i>
Next state

501

502 Each state supports a number of possible events. Expected events are processed by taking the
 503 prescribed action and transitioning to the next state. Unexpected protocol messages will result in a fault
 504 message, with a standard fault code such as Invalid State or Inconsistent Internal State. Events that may
 505 not occur in a given state are labeled as N/A.

506 Notes:

- 507 1. Transitions with a "N/A" as their action are inexpressible. A TM should view these transitions as
 508 serious internal consistency issues, and probably fatal.
- 509 2. The "Internal events" shown are those events, created either within a TM itself or on its local
 510 system, that cause state changes and/or trigger the sending of a protocol message.

511 **9.1 Completion Protocol**

512

Completion Protocol (Coordinator View)			
Inbound Events	States		
	None	Active	Completing
Commit	<i>Unknown Transaction</i> None	<i>Initiate user commit</i> Completing	<i>Ignore</i> Completing
Rollback	<i>Unknown Transaction</i> None	<i>Initiate user rollback, send aborted</i> None	<i>Invalid State</i> Completing
Internal Events			
Commit Decision	N/A	N/A	<i>Send committed</i> None
Abort Decision	N/A	<i>Send aborted</i> None	<i>Send aborted</i> None

513

514 **9.2 2PC Protocol**

515 These tables present the view of a coordinator or participant with respect to a single partner. A
 516 coordinator with multiple participants can be understood as a collection of independent coordinator state
 517 machines, each with its own state.

518

Atomic Transaction 2PC Protocol (Coordinator View)							
Inbound Events	States						
	None	Active	Preparing	Prepared	PreparedSuccess	Committing	Aborting
Prepared	<i>Durable: Send Rollback</i> <i>Volatile: Unknown Transaction</i> None	<i>Invalid State</i> Aborting	<i>Record Vote</i> Prepared	<i>Ignore</i> Prepared	<i>Ignore</i> PreparedSuccess	<i>Resend Commit</i> Committing	<i>Resend Rollback</i> Aborting
ReadOnly	<i>Ignore</i> None	<i>Forget</i> None	<i>Forget</i> None	<i>Inconsistent Internal State</i> Prepared	<i>Inconsistent Internal State</i> PreparedSuccess	<i>Inconsistent Internal State</i> Committing	<i>Forget</i> None
Aborted	<i>Ignore</i> None	<i>Forget</i> None	<i>Forget</i> None	<i>Inconsistent Internal State</i> Prepared	<i>Inconsistent Internal State</i> PreparedSuccess	<i>Inconsistent Internal State</i> Committing	<i>Forget</i> None
Committed	<i>Ignore</i> None	<i>Invalid State</i> Aborting	<i>Invalid State</i> Aborting	<i>Inconsistent Internal State</i> Prepared	<i>Inconsistent Internal State</i> PreparedSuccess	<i>Forget</i> None	<i>Inconsistent Internal State</i> Aborting
Internal Events							
User Commit	<i>N/A</i>	<i>Send Prepare</i> Preparing	<i>N/A</i>	<i>N/A</i>	<i>N/A</i>	<i>N/A</i>	<i>N/A</i>
User Rollback	<i>N/A</i>	<i>Send Rollback</i> Aborting	<i>N/A</i>	<i>N/A</i>	<i>N/A</i>	<i>N/A</i>	<i>N/A</i>
Expires Times Out	<i>N/A</i>	<i>Send Rollback</i> Aborting	<i>Send Rollback</i> Aborting	<i>Send Rollback</i> Aborting	<i>Ignore</i> PreparedSuccess	<i>Ignore</i> Committing	<i>Ignore</i> Aborting
Comms Times Out	<i>N/A</i>	<i>N/A</i>	<i>Resend Prepare</i> Preparing	<i>N/A</i>	<i>N/A</i>	<i>Resend Commit</i> Committing	<i>N/A</i>
Commit Decision	<i>N/A</i>	<i>N/A</i>	<i>N/A</i>	<i>Record Outcome</i> PreparedSuccess	<i>N/A</i>	<i>N/A</i>	<i>N/A</i>
Rollback Decision	<i>N/A</i>	<i>Send Rollback</i> Aborting	<i>Send Rollback</i> Aborting	<i>Send Rollback</i> Aborting	<i>N/A</i>	<i>N/A</i>	<i>N/A</i>
Write Done	<i>N/A</i>	<i>N/A</i>	<i>N/A</i>	<i>N/A</i>	<i>Send Commit</i> Committing	<i>N/A</i>	<i>N/A</i>
Write Failed	<i>N/A</i>	<i>N/A</i>	<i>N/A</i>	<i>N/A</i>	<i>Send Rollback</i> Aborting	<i>N/A</i>	<i>N/A</i>
Participant Abandoned	<i>N/A</i>	<i>N/A</i>	<i>N/A</i>	<i>N/A</i>	<i>N/A</i>	<i>Durable: N/A</i> <i>Volatile: None</i>	<i>None</i>

519

520 “Forget” implies that the subordinate’s participation is removed from the coordinator (if necessary), and
 521 otherwise the message is ignored

Atomic Transaction 2PC Protocol (Participant View)						
Inbound Events	States					
	None	Active	Preparing	Prepared	PreparedSuccess	Committing
Prepare	<i>Send Aborted</i> None	<i>Gather Vote</i> <i>Decision</i> Preparing	<i>Ignore</i> Preparing	<i>Ignore</i> Prepared	<i>Resend Prepared</i> PreparedSuccess	<i>Ignore</i> Committing
Commit	<i>Send Committed</i> None	<i>Invalid State</i> None	<i>Invalid State</i> None	<i>Invalid State</i> None	<i>Initiate Commit Decision</i> Committing	<i>Ignore</i> Committing
Rollback	<i>Send Aborted</i> None	<i>Initiate Rollback and Send Aborted</i> None	<i>Initiate Rollback and Send Aborted</i> None	<i>Initiate Rollback and Send Aborted</i> None	<i>Initiate Rollback and Send Aborted</i> None	<i>Inconsistent Internal State</i> Committing
Internal Events						
Expires Times Out	<i>N/A</i>	<i>Initiate Rollback and Send Aborted</i> None	<i>Initiate Rollback and Send Aborted</i> None	<i>Ignore</i> Prepared	<i>Ignore</i> PreparedSuccess	<i>Ignore</i> Committing
Comms Times Out	<i>N/A</i>	<i>N/A</i>	<i>N/A</i>	<i>N/A</i>	<i>Resend Prepared</i> PreparedSuccess	<i>N/A</i>
Commit Decision	<i>N/A</i>	<i>N/A</i>	<i>Record Commit</i> Prepared	<i>N/A</i>	<i>N/A</i>	<i>Send Committed</i> None
Rollback Decision	<i>N/A</i>	<i>Send Aborted</i> None	<i>Send Aborted</i> None	<i>N/A</i>	<i>N/A</i>	<i>N/A</i>
Write Done	<i>N/A</i>	<i>N/A</i>	<i>N/A</i>	<i>Send Prepared</i> PreparedSuccess	<i>N/A</i>	<i>N/A</i>
Write Failed	<i>N/A</i>	<i>N/A</i>	<i>N/A</i>	<i>Initiate Rollback and Send Aborted</i> None	<i>N/A</i>	<i>N/A</i>
ReadOnly Decision	<i>N/A</i>	<i>Send ReadOnly</i> None	<i>Send ReadOnly</i> None	<i>N/A</i>	<i>N/A</i>	<i>N/A</i>

522

523

A. Acknowledgements

524 This document is based on initial contributions to the OASIS WS-TX Technical Committee by the
525 following authors: Luis Felipe Cabrera (Microsoft), George Copeland (Microsoft), Max Feingold
526 (Microsoft), Robert W Freund (Hitachi), Tom Freund (IBM), Sean Joyce (IONA), Johannes Klein
527 (Microsoft), David Langworthy (Microsoft), Mark Little (Arjuna Technologies), Frank Leymann (IBM), Eric
528 Newcomer (IONA), David Orchard (BEA Systems), Ian Robinson (IBM), Tony Storey (IBM), Satish Thatte
529 (Microsoft).

530

531 The following individuals have provided invaluable input into the initial contribution: Francisco Curbera
532 (IBM), Doug Davis (IBM), Gert Drapers (Microsoft), Don Ferguson (IBM), Kirill Gavrylyuk (Microsoft), Dan
533 House (IBM), Oisin Hurley (IONA), Thomas Mikalsen (IBM), Jagan Peri (Microsoft), John Shewchuk
534 (Microsoft), Stefan Tai (IBM).

535

536 The following individuals have participated in the creation of this specification and are gratefully
537 acknowledged:

538 **Participants:**

539 Martin Chapman, Oracle
540 Kevin Conner, JBoss Inc.
541 Paul Cotton, Microsoft Corporation
542 Doug Davis, IBM
543 Colleen Evans, Microsoft Corporation
544 Max Feingold, Microsoft Corporation
545 Thomas Freund, IBM
546 Robert Freund, Hitachi, Ltd.
547 Peter Furniss, Erebor Ltd.
548 Marc Goodner, Microsoft Corporation
549 Alastair Green, Choreology Ltd.
550 Daniel House, IBM
551 Ram Jeyaraman, Microsoft Corporation
552 Paul Knight, Nortel Networks Limited
553 Mark Little, JBoss Inc.
554 Jonathan Marsh, Microsoft Corporation
555 Monica Martin, Sun Microsystems
556 Joseph Fialli, Sun Microsystems
557 Eric Newcomer, IONA Technologies
558 Eisaku Nishiyama, Hitachi, Ltd.
559 Alain Regnier, Ricoh Company, Ltd.
560 Ian Robinson, IBM
561 Tom Rutt, Fujitsu Limited
562 Andrew Wilkinson, IBM

B. Revision History

Revision	yy-mm-dd	Editor	Changes Made
01	05-11-22	Mark Little Andrew Wilkinson	Initial Working Draft
02	06-02-12	Mark Little	Updated for issue i017
03	06-03-02	Andrew Wilkinson	Updated for issue i015
04	06-03-10	Andrew Wilkinson	Updated for issue i009
cd-01	06-03-15	Andrew Wilkinson	Updates to produce CD-01
05	06-05-23	Mark Little	Updates for i023, i026, i027, i028, i030
06	06-06-01	Andrew Wilkinson	Updates for i039, i043, i045, i052, i053, i055
cd-02	06-06-13	Andrew Wilkinson	Updates to produce CD-02
07	06-07-13	Mark Little	Editorial changes.
08	06-07-24	Andrew Wilkinson	Updates for i036, i037. Update namespace to 2006/06
09	06-08-18	Mark Little Andrew Wilkinson	Updates for i038, i041, i047, i049, i050, i056, i057, i062, i083, i084.
10	06-08-25	Andrew Wilkinson	Updates for i061, i065, i078, i080, i081, i089
11	06-08-30	Andrew Wilkinson	Editorial changes Updates for i090, i091
cd-03	06-08-30	Andrew Wilkinson	Updates to produce CD-03