



Web Services Atomic Transaction (WS-AtomicTransaction) 1.1

Committee Draft 01, March 15, 2006

Document Identifier:

wstx-wsat-1.1-spec-cd-01

Location:

<http://docs.oasis-open.org/ws-tx/wstx-wsat-1.1-spec-cd-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 "committee 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 does 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	Note on terminology.....	4
1.1	Composable Architecture.....	4
1.2	Namespace.....	4
1.2.1	Prefix Namespace.....	4
1.3	XSD and WSDL Files.....	4
1.4	AT Protocol Elements.....	5
1.5	Normative References.....	5
1.6	Non-normative References.....	5
2	Introduction.....	7
3	Atomic Transaction Context.....	8
4	Atomic Transaction Protocols.....	9
4.1	Preconditions.....	9
4.2	Completion Protocol.....	9
4.3	Two-Phase Commit Protocol.....	10
4.3.1	Volatile Two-Phase Commit Protocol.....	10
4.3.2	Durable Two-Phase Commit Protocol.....	11
4.3.3	2PC Diagram and Notifications.....	11
5	AT Policy Assertion.....	13
5.1	Assertion Model.....	13
5.2	Normative Outline.....	13
5.3	Assertion Attachment.....	14
5.4	Assertion Example.....	14
6	Transaction Faults.....	16
6.1	InconsistentInternalState.....	17
7	Security Model.....	18
8	Security Considerations.....	20
9	Use of WS-Addressing Headers.....	22
10	State Tables.....	23
	Appendix A. Acknowledgements.....	26
	Appendix B. Revision History.....	27

1 Note on terminology

The keywords "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 <http://example.org> and <http://example.com> represents some application-dependent or context-dependent URI as defined in RFC 2396 [URI].

1.1 Composable Architecture

By using the SOAP [SOAP] 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 Namespace

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

```
http://docs.oasis-open.org/ws-tx/wsat/2006/03
```

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

1.2.1 Prefix Namespace

Prefix	Namespace
S	http://www.w3.org/2003/05/soap-envelope
wscor	http://docs.oasis-open.org/ws-tx/wscor/2006/03
wsat	http://docs.oasis-open.org/ws-tx/wsat/2006/03

If an action URI is used then the action URI MUST consist of the wsat namespace URI concatenated with the "/" character and the element name. For example:

```
http://docs.oasis-open.org/ws-tx/wsat/2006/03/Commit
```

1.3 XSD and WSDL Files

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

<http://docs.oasis-open.org/ws-tx/wsat/2006/03/wsat.xsd>

<http://docs.oasis-open.org/ws-tx/wsat/2006/03/wsat.wsdl>

Soap bindings for the WSDL documents defined in this specification MUST use "document" for the *style* attribute.

28 **1.4 AT Protocol Elements**

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

33 **1.5 Normative References**

34 **1.6 Non-normative References**

35 **[KEYWORDS]**

36 S. Bradner, "Key words for use in RFCs to Indicate Requirement Levels," [RFC 2119](#), Harvard
37 University, March 1997

38 **[SOAP]**

39 W3C Note, "[SOAP: Simple Object Access Protocol 1.1](#)," 08 May 2000

40 **[URI]**

41 T. Berners-Lee, R. Fielding, L. Masinter, "Uniform Resource Identifiers (URI): Generic Syntax,"
42 [RFC 2396](#), MIT/LCS, U.C. Irvine, Xerox Corporation, August 1998

43 **[XML-ns]**

44 W3C Recommendation, "[Namespaces in XML](#)," 14 January 1999

45 **[XML-Schema1]**

46 W3C Recommendation, "[XML Schema Part 1: Structures](#)," 2 May 2001

47 **[XML-Schema2]**

48 W3C Recommendation, "[XML Schema Part 2: Datatypes](#)," 2 May 2001

49 **[WSCOOR]**

50 [Web Services Coordination \(WS-Coordination\) 1.1](#), OASIS, March 2006

51 **[WSADDR]**

52 [Web Services Addressing \(WS-Addressing\)](#), Microsoft, IBM, Sun, BEA Systems, SAP, Sun,
53 August 2004

54 **[WSPOLICY]**

55 [Web Services Policy Framework \(WS-Policy\)](#), VeriSign, Microsoft, Sonic Software, IBM, BEA
56 Systems, SAP, September 2004

57 **[WSPOLICYATTACH]**

58 [Web Services Policy Attachment \(WS-PolicyAttachment\)](#), VeriSign, Microsoft, Sonic Software,
59 IBM, BEA Systems, SAP, September 2004

60 **[WSDL]**

61 Web Services Description Language (WSDL) 1.1

62 "<http://www.w3.org/TR/2001/NOTE-wsdl-20010315>"

- 63 **[WSec]**
64 OASIS Standard 200401, March 2004, "[Web Services Security: SOAP Message Security 1.0](#)
65 (WS-Security 2004)"
- 66 **[WSecPolicy]**
67 [Web Services Security Policy Language \(WS-SecurityPolicy\)](#), Microsoft, VeriSign, IBM, RSA
68 Security, July 2005
- 69 **[WSecConv]**
70 [Web Services Secure Conversation Language \(WS-SecureConversation\)](#), OpenNetwork, Layer7,
71 Netegrity, Microsoft, Reactivity, IBM, VeriSign, BEA Systems, Oblix, RSA Security, Ping Identity,
72 Westbridge, Computer Associates, February 2005
- 73 **[WSTrust]**
74 [Web Services Trust Language \(WS-Trust\)](#), OpenNetwork, Layer7, Netegrity, Microsoft,
75 Reactivity, VeriSign, IBM, BEA Systems, Oblix, RSA Security, Ping Identity, Westbridge,
76 Computer Associates, February 2005.

77 2 Introduction

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

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

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

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

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

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

108 3 Atomic Transaction Context

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

112 The Atomic Transaction coordination context must flow on all application messages involved with the
113 transaction.

114 Atomic Transaction adds the following semantics to the CreateCoordinationContext operation on the
115 activation service.

- 116 • If the request includes the CurrentContext element, the target coordinator is interposed as a
117 subordinate to the coordinator stipulated inside the CurrentContext element.
- 118 • If the request does not include a CurrentContext element, the target coordinator creates a new
119 transaction and acts as the root.

120 A coordination context may have an Expires attribute. This attribute specifies the earliest point in time at
121 which a transaction may be terminated solely due to its length of operation. From that point forward, the
122 transaction manager may elect to unilaterally roll back the transaction, so long as it has not transmitted a
123 Commit or a Prepared notification.

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

125 `http://docs.oasis-open.org/ws-tx/wsat/2006/03`

126 4 Atomic Transaction Protocols

127 This specification defines the following protocols for atomic transactions.

- 128 • **Completion:** The completion protocol initiates commitment processing. Based on each
129 protocol's registered participants, the coordinator begins with Volatile 2PC then proceeds through
130 Durable 2PC. The final result is signaled to the initiator.
- 131 • **Two-Phase Commit (2PC):** The 2PC protocol coordinates registered participants to reach a
132 commit or abort decision, and ensures that all participants are informed of the final result. The
133 2PC protocol has two variants:
 - 134 ○ **Volatile 2PC:** Participants managing volatile resources such as a cache should
135 register for this protocol.
 - 136 ○ **Durable 2PC:** Participants managing durable resources such as a database should
137 register for this protocol.

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

139 4.1 Preconditions

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

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

146 4.2 Completion Protocol

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

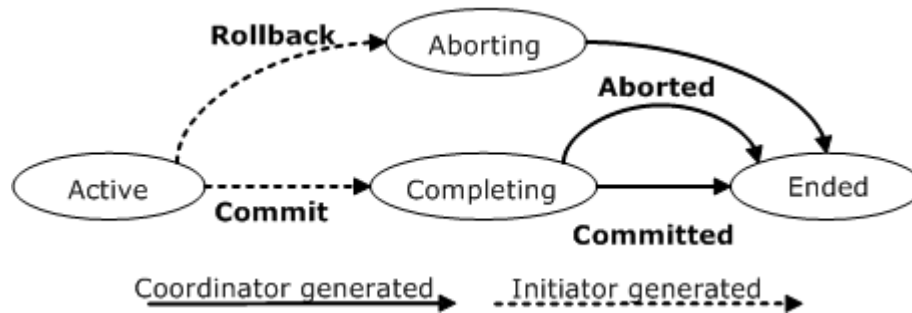
149 An initiator registers for this protocol using the following protocol identifier:

150 `http://docs.oasis-open.org/ws-tx/wsac/2006/03/Completion`

151

152 The diagram below illustrates the protocol abstractly:

153



154

155

156 The coordinator accepts:

157 Commit

158 Upon receipt of this notification, the coordinator knows that the participant has completed
159 application processing and that it should attempt to commit the transaction.

160 Rollback

161 Upon receipt of this notification, the coordinator knows that the participant has terminated
162 application processing and that it should abort the transaction.

163 The initiator accepts:

164 Committed

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

167 Aborted

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

170 Conforming implementations must implement Completion.

171 4.3 Two-Phase Commit Protocol

172 The Two-Phase Commit (2PC) protocol is a Coordination protocol that defines how multiple participants
173 reach agreement on the outcome of an atomic transaction. The 2PC protocol has two variants: Durable
174 2PC and Volatile 2PC.

175 4.3.1 Volatile Two-Phase Commit Protocol

176 Upon receiving a Commit notification in the completion protocol, the root coordinator begins the prepare
177 phase of all participants registered for the Volatile 2PC protocol. All participants registered for this
178 protocol must respond before a Prepare is issued to a participant registered for Durable 2PC. Further
179 participants may register with the coordinator until the coordinator issues a Prepare to any durable
180 participant. A volatile recipient is not guaranteed to receive a notification of the transaction's outcome.

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

182

<http://docs.oasis-open.org/ws-tx/wsac/2006/03/Volatile2PC>

183 **4.3.2 Durable Two-Phase Commit Protocol**

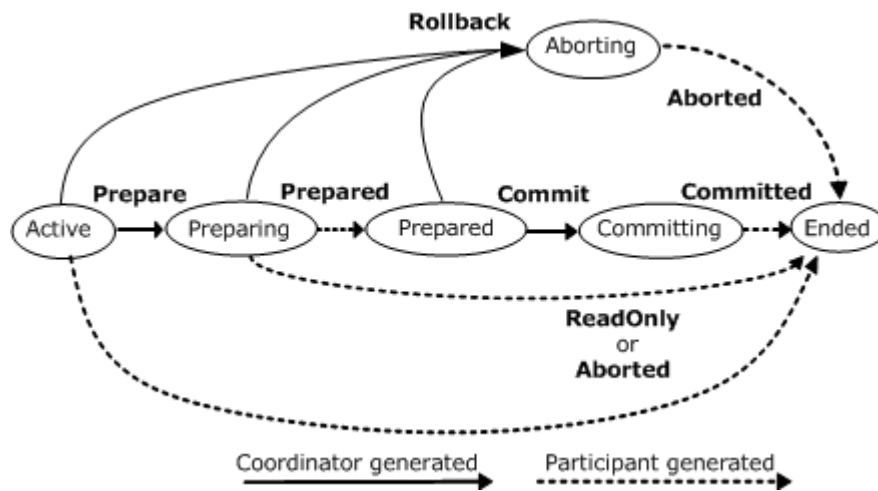
184 After receiving a Commit notification in the completion protocol and upon successfully completing the
185 prepare phase for Volatile 2PC participants, the root coordinator begins the Prepare phase for Durable
186 2PC participants. All participants registered for this protocol must respond Prepared or ReadOnly before
187 a Commit notification is issued to a participant registered for either protocol.

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

189 <http://docs.oasis-open.org/ws-tx/wsac/2006/03/Durable2PC>

190 **4.3.3 2PC Diagram and Notifications**

191 The diagram below illustrates the protocol abstractly:



194 The participant accepts:

195 Prepare

196 Upon receipt of this notification, the participant knows to enter phase 1 and vote on the outcome
197 of the transaction. If the participant does not know of the transaction, it must vote to abort. If the
198 participant has already voted, it should resend the same vote.

199 Rollback

200 Upon receipt of this notification, the participant knows to abort, and forget, the transaction. This
201 notification can be sent in either phase 1 or phase 2. Once sent, the coordinator may forget all
202 knowledge of this transaction.

203 Commit

204 Upon receipt of this notification, the participant knows to commit the transaction. This notification
205 can only be sent after phase 1 and if the participant voted to commit. If the participant does not
206 know of the transaction, it must send a Committed notification to the coordinator.

207 The coordinator accepts:

208 Prepared

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

- 211 ReadOnly
- 212 Upon receipt of this notification, the coordinator knows the participant votes to commit the
213 transaction, and has forgotten the transaction. The participant does not wish to participate in
214 phase 2.
- 215 Aborted
- 216 Upon receipt of this notification, the coordinator knows the participant has aborted, and forgotten,
217 the transaction.
- 218 Committed
- 219 Upon receipt of this notification, the coordinator knows the participant has committed the
220 transaction. That participant may be safely forgotten.
- 221 Replay
- 222 Upon receipt of this notification, the coordinator may assume the participant has suffered a
223 recoverable failure. It should resend the last appropriate protocol notification.
- 224 Conforming implementations MUST implement the 2PC protocol.

225

5 AT Policy Assertion

226 WS-Policy Framework [\[WS-Policy\]](#) and WS-Policy Attachment [\[WS-PolicyAttachment\]](#) collectively define
227 a framework, model and grammar for expressing the capabilities, requirements, and general
228 characteristics of entities in an XML Web services-based system. To enable a web service to describe
229 transactional capabilities and requirements of a service and its operations, this specification defines a pair
230 of Atomic Transaction policy assertions that leverage the WS-Policy framework.

5.1 Assertion Model

232 The AT policy assertions are provided by a web service to qualify the transactional processing of
233 messages associated with the particular operation to which the assertions are scoped. The AT policy
234 assertions indicate:

- 235 1. whether a requester MAY, MUST or SHOULD NOT include an AtomicTransaction
236 CoordinationContext flowed with the message.
- 237 2. the capability of the target service to process the message under an atomic transaction
238 regardless of whether the requester supplies an AtomicTransaction CoordinationContext.

239 The AT policy assertions are semantically independent of one another, and may be used together or in
240 isolation.

5.2 Normative Outline

242 The normative outlines for the AT policy assertions are:

```
243 <wsat:ATAssertion [wsp:Optional="true"]? ... >  
244 ...  
245 </wsat:ATAssertion>
```

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

247 /wsat:ATAssertion

248 A policy assertion that specifies that an atomic transaction MUST be flowed inside a requester's
249 message. From the perspective of the requester, the target service that processes the transaction MUST
250 behave as if it had participated in the transaction. The transaction MUST be represented as a SOAP
251 header in CoordinationContext format, as defined in WS-Coordination [\[WS-Coordination\]](#).

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

253 Per WS-Policy [\[WS-Policy\]](#), this is compact notation for two policy alternatives, one with and one without
254 the assertion. Presence of both policy alternatives indicates that the behavior indicated by the assertion is
255 optional, such that an atomic transaction MAY be flowed inside a requester's message. The absence of
256 the assertion is interpreted to mean that a transaction SHOULD NOT be flowed inside a requester's
257 message.

```
258 <wsat:ATAlwaysCapability ... />
```

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

260 /wsat:ATAlwaysCapability

261 A policy assertion that specifies a capability of the target service indicating that a requester's message
262 will be processed transactionally regardless of whether the requester supplies an AtomicTransaction
263 CoordinationContext. If an AtomicTransaction context is provided by the requester, it will be used.
264 Otherwise the processing of the message will be within a transaction implicitly started and ended by the
265 target service's environment as part of the processing of that message.

266 5.3 Assertion Attachment

267 Because the AT policy assertions indicate atomic transaction behavior for a single operation, the
268 assertions have Operation Policy Subject [[WS-PolicyAttachment](#)].

269 WS-PolicyAttachment defines two WSDL [[WSDL 1.1](#)] policy attachment points with Operation Policy
270 Subject:

- 271 • wsdl:portType/wsdl:operation – A policy expression containing the AT policy assertion MUST
272 NOT be attached to a wsdl:portType; the AT policy assertions specify a concrete behavior
273 whereas the wsdl:portType is an abstract construct.
- 274 • wsdl:binding/wsdl:operation – A policy expression containing the AT policy assertions SHOULD
275 be attached to a wsdl:binding.

276 5.4 Assertion Example

277 An example use of the AT policy assertion follows:

```
278 (01) <wsdl:definitions  
279 (02)     targetNamespace="bank.example.com"  
280 (03)     xmlns:tns="bank.example.com"  
281 (04)     xmlns:wSDL="http://schemas.xmlsoap.org/wSDL/"  
282 (05)     xmlns:wsp="http://schemas.xmlsoap.org/ws/2004/09/policy"  
283 (06)     xmlns:wsat="http://docs.oasis-open.org/ws-tx/wsat/2006/03"  
284 (07)     xmlns:wssu="http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-  
285 wssecurity-utility-1.0.xsd" >  
286 (08)  
287 (09)     <wsp:Policy wsu:Id="TransactedPolicy1" >  
288 (10)         <wsat:ATAssertion wsp:optional="true" />  
289 (11)         <!-- omitted assertions -->  
290 (12)     </wsp:Policy>  
291 (13)     <wsp:Policy wsu:Id="TransactedPolicy2" >  
292 (14)         <wsat:ATAlwaysCapability />  
293 (15)         <!-- omitted assertions -->  
294 (16)     </wsp:Policy>  
295 (17)     <!-- omitted elements -->  
296 (18)     <wsdl:binding name="BankBinding" type="tns:BankPortType" >  
297 (19)         <!-- omitted elements -->
```

```

298 (20)      <wsdl:operation name="QueryBalance" >
299 (21)      <wsp:PolicyReference URI="#TransactedPolicy2"
300 wsdl:required="true" />
301 (22)      <!-- omitted elements -->
302 (23)      </wsdl:operation>
303 (24)      <wsdl:operation name="TransferFunds" >
304 (25)      <wsp:PolicyReference URI="#TransactedPolicy1"
305 wsdl:required="true" />
306 (26)      <!-- omitted elements -->
307 (27)      </wsdl:operation>
308 (28)      </wsdl:binding>
309 (29)      </wsdl:definitions>

```

310

311 Lines (9-12) are a policy expression that includes an AT policy assertion (Line 10) to indicate that an
312 atomic transaction in WS-Coordination [\[WS-Coordination\]](#) format MAY be used.

313 Lines (13-16) are a policy expression that includes an AT policy assertion (Line 14) to indicate that a
314 capability of the target service is that it will process messages in a transaction regardless of whether any
315 AtomicTransaction CoordinationContext is sent by the requester.

316 Lines (20-23) are a WSDL [\[WSDL 1.1\]](#) binding. Line (21) indicates that the policy in Lines (13-16) applies
317 to this binding, specifically indicating that QueryBalance messages are processed in an atomic
318 transaction regardless of whether a requester provides an AtomicTransaction CoordinationContext.

319 Lines (24-27) are a WSDL [\[WSDL 1.1\]](#) binding. Line (25) indicates that the policy in Lines (9-12) applies
320 to this binding, specifically indicating that an atomic transaction MAY flow inside messages.

321

6 Transaction Faults

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

323

```
http://docs.oasis-open.org/ws-tx/wsat/2006/03/fault
```

324 The faults defined in this section are generated if the condition stated in the preamble is met. Faults are
325 targeted at a destination endpoint according to the fault handling rules defined in [WSADDR].

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

327 [Code] The fault code.

328 [Subcode] The fault subcode.

329 [Reason] The English language reason element.

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

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

333

SOAP Version	Sender	Receiver
SOAP 1.2	S:Sender	S:Receiver

334

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

336

```

<S:Envelope>
337 <S:Header>
338   <wsa:Action>
339     http://docs.oasis-open.org/ws-tx/wsat/2006/03/fault
340   </wsa:Action>
341   <!-- Headers elided for clarity. -->
342 </S:Header>
343 <S:Body>
344   <S:Fault>
345     <S:Code>
346       <S:Value>[Code]</S:Value>
347       <S:Subcode>
348         <S:Value>[Subcode]</S:Value>
349       </S:Subcode>
350     </S:Code>
351     <S:Reason>
352       <S:Text xml:lang="en">[Reason]</S:Text>
353     </S:Reason>
354     <S:Detail>
355       [Detail]
356       ...
357     </S:Detail>
358   </S:Fault>
359 </S:Body>
360 </S:Envelope>

```

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

362

```
<S11:Envelope>
```



```
363 <S11:Body>
364 <S11:Fault>
365 <faultcode>[Subcode]</faultcode>
366 <faultstring xml:lang="en">[Reason]</faultstring>
367 </S11:Fault>
368 </S11:Body>
369 </S11:Envelope>
```

370 **6.1 InconsistentInternalState**

371 This fault is sent by a participant to indicate that it cannot fulfill its obligations. This indicates a global
372 consistency failure and is an unrecoverable condition.

373 Properties:

374 **[Code]** Sender

375 **[Subcode]** wsat:InconsistentInternalState

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

377 **[Detail]** unspecified

378

7 Security Model

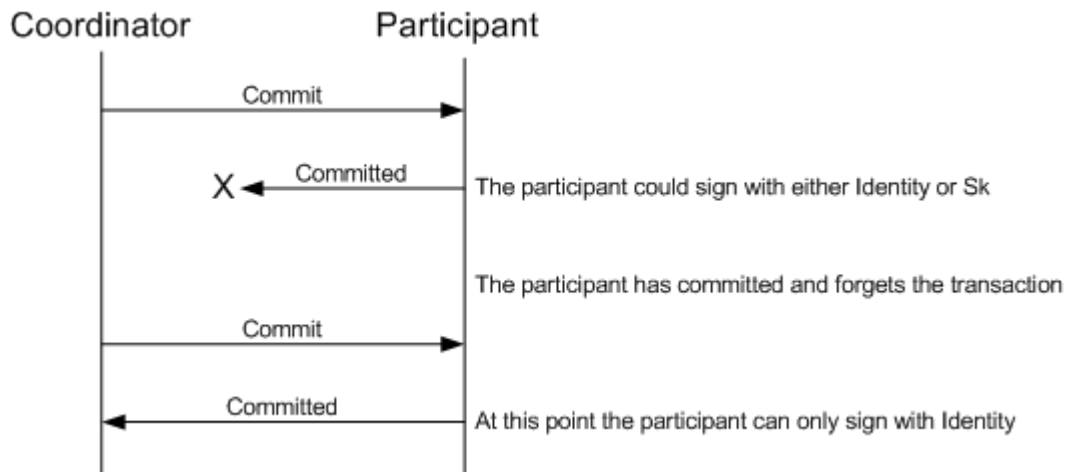
379 The security model for atomic transactions builds on the model defined in WS-Coordination [WSCOOR].
380 That is, services have policies specifying their requirements and requestors provide claims (either implicit
381 or explicit) and the requisite proof of those claims. Coordination context creation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

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

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

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

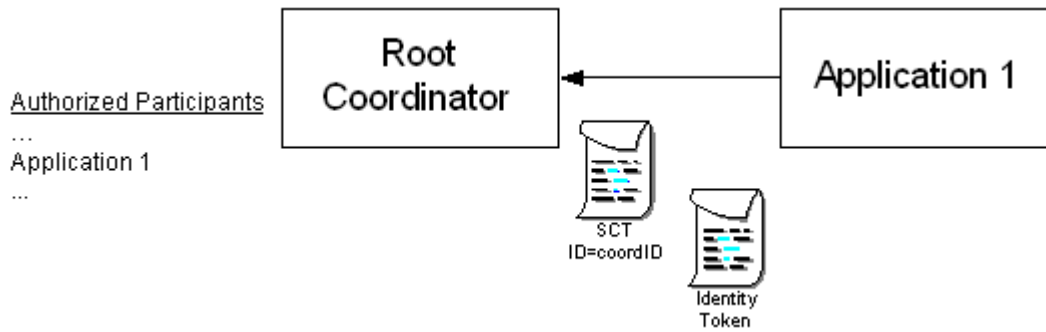
393 The participant can only prove its identity to the coordinator when it indicates that the specified
394 transaction is not in its log and assumed committed. This is illustrated in the figure below:



395

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

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



403

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

407

8 Security Considerations

408 It is strongly RECOMMENDED that the communication between services be secured using the
409 mechanisms described in WS-Security [WSSec]. In order to properly secure messages, the body and all
410 relevant headers need to be included in the signature. Specifically, the
411 <wscoor:CoordinationContext> header needs to be signed with the body and other key message
412 headers in order to "bind" the two together.

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

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

- 420 • Attaching a nonce to each message and using it in a derived key function with the shared secret
- 421 • Using a derived key sequence and switch "generations"
- 422 • Closing and re-establishing a security context (not possible for delegated keys)
- 423 • Exchanging new secrets between the parties (not possible for delegated keys)

424 It should be noted that the mechanisms listed above are independent of the SCT and secret returned
425 when the coordination context is created. That is, the keys used to secure the channel may be
426 independent of the key used to prove the right to register with the activity.

427 The security context MAY be re-established using the mechanisms described in WS-Trust [WSTrust] and
428 WS-SecureConversation [WSSecConv]. Similarly, secrets can be exchanged using the mechanisms
429 described in WS-Trust. Note, however, that the current shared secret SHOULD NOT be used to encrypt
430 the new shared secret. Derived keys, the preferred solution from this list, can be specified using the
431 mechanisms described in WS-SecureConversation.

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

- 434 • **Message alteration** – Alteration is prevented by including signatures of the message information
435 using WS-Security [WSSec].
- 436 • **Message disclosure** – Confidentiality is preserved by encrypting sensitive data using WS-
437 Security.
- 438 • **Key integrity** – Key integrity is maintained by using the strongest algorithms possible (by
439 comparing secured policies – see WS-Policy [WSPOLICY] and WS-SecurityPolicy
440 [WSSecPolicy]).
- 441 • **Authentication** – Authentication is established using the mechanisms described in WS-Security
442 and WS-Trust [WSTrust]. Each message is authenticated using the mechanisms described in
443 WS-Security [WSSec].
- 444 • **Accountability** – Accountability is a function of the type of and string of the key and algorithms
445 being used. In many cases, a strong symmetric key provides sufficient accountability. However,
446 in some environments, strong PKI signatures are required.

- 447
- 448
- 449
- 450
- 451
- **Availability** – Many services are subject to a variety of availability attacks. Replay is a common attack and it is RECOMMENDED that this be addressed as described in the next bullet. Other attacks, such as network-level denial of service attacks are harder to avoid and are outside the scope of this specification. That said, care should be taken to ensure that minimal processing be performed prior to any authenticating sequences.
- 452
- **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.
- 453
- 454
- 455

456

9 Use of WS-Addressing Headers

457 The messages defined in WS-AtomicTransaction can be classified into two types:

458 • Notification messages: **Commit, Rollback, Committed, Aborted, Prepare,**
459 **Prepared, ReadOnly** and **Replay**.

460 • Fault messages

461 Notification messages follow the standard "one way" pattern as defined in WS-Addressing. There are two
462 types of notification messages:

463 • A notification message is a terminal message when it indicates the end of a
464 coordinator/participant relationship. **Committed, Aborted** and **ReadOnly** are
465 terminal messages.

466 • A notification message is a non-terminal message when it does not indicate the end
467 of a coordinator/participant relationship. **Commit, Rollback, Prepare, Prepared**
468 and **Replay** are non-terminal messages.

469 The following statements define addressing interoperability requirements for the WS-AtomicTransaction
470 message types:

471 Non-terminal notification messages

472 • MUST include a wsa:ReplyTo header

473 Terminal notification messages

474 • SHOULD NOT include a wsa:ReplyTo header

475 Fault messages

476 • MUST include a wsa:RelatesTo header, specifying the MessageID from the Notification
477 message that generated the fault condition.

478

479 Notification messages are addressed by both coordinators and participants using the Endpoint
480 References initially obtained during the Register-RegisterResponse exchange. If a wsa:ReplyTo header
481 is present in a notification message it MAY be used by the recipient, for example in cases where a
482 Coordinator or Participant has forgotten a transaction that is completed and needs to respond to a resent
483 protocol message. Permanent loss of connectivity between a coordinator and a participant in an in-doubt
484 state can result in data corruption.

485 If a wsa:FaultTo header is present on a message that generates a fault condition, then it MUST be used
486 by the recipient as the destination for any fault. Otherwise, fault messages MAY be addressed by both
487 coordinators and participants using the Endpoint References initially obtained during the Register-
488 RegisterResponse exchange.

489 All messages are delivered using connections initiated by the sender. Endpoint References MUST
490 contain physical addresses and MUST NOT use the well-known "anonymous" endpoint defined in WS-
491 Addressing.

492

10 State Tables

493

494

495

496

The following state tables specify the behavior of coordinators and participants when presented with protocol messages or internal events. These tables present the view of a coordinator or participant with respect to a single partner. A coordinator with multiple participants can be understood as a collection of independent coordinator state machines.

497

Each cell in the tables uses the following convention:

498

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

499

500

501

502

503

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

Atomic Transaction 2PC protocol (Coordinator View)							
Inbound Events	States						
	None	Active	Preparing	Prepared	PreparedSuccess	Committing	Aborting
Register	Invalid State None	Send RegisterResponse Active	Durable: Invalid State Aborting Volatile: Send RegisterResponse Active	N/A	Invalid State PreparedSuccess	Invalid State Committing	Invalid State Aborting
Prepared	Durable: Send Rollback Volatile: Invalid State None	Invalid State Aborting	Record Vote Preparing	N/A	Ignore PreparedSuccess	Resend Commit Committing	Resend Rollback, and Forget Aborting
ReadOnly	Ignore None	Forget Active	Forget Preparing	N/A	Invalid State PreparedSuccess	Invalid State Committing	Forget Aborting
Aborted	Ignore None	Forget Aborting	Forget Aborting	N/A	Invalid State PreparedSuccess	Invalid State Committing	Forget Aborting
Committed	Ignore None	Invalid State Aborting	Invalid State Aborting	N/A	Invalid State PreparedSuccess	Forget Committing	Invalid State Aborting
Replay	Durable: Send Rollback Volatile: Invalid State None	Send Rollback Aborting	Send Rollback Aborting	N/A	Ignore PreparedSuccess	Send Commit Committing	Send Rollback Aborting
Internal Events							
User Commit	Return Aborted None	Send Prepare Preparing	Ignore Preparing	N/A	Ignore Prepared Success	Return Committed Committing	Return Aborted Aborting
User Rollback	Return Aborted None	Send Rollback Aborting	Send Rollback Aborting	N/A	Invalid State PreparedSuccess	Invalid State Committing	Return Aborted Aborting
Expires Times Out	N/A	Send Rollback Aborting	Send Rollback Aborting	N/A	Ignore PreparedSuccess	Ignore Committing	Ignore Aborting
Comms Times Out	N/A	N/A	Resend Prepare Preparing	N/A	N/A	Resend Commit Committing	N/A
Commit Decision	N/A	N/A	Record Outcome Prepared Success	N/A	N/A	N/A	N/A
Write Done	N/A	N/A	N/A	N/A	Send Commit Committing	N/A	N/A
Write Failed	N/A	N/A	N/A	N/A	Send Rollback Aborting	N/A	N/A
All Forgotten	N/A	Active	None	N/A	N/A	None	None

504

505 Notes:

506 1. Transitions with a “N/A” as their action are inexpressible. A TM should view these transitions as
507 serious internal consistency issues, and probably fatal.

508 2. Internal events are those that are created either within a TM itself, or on its local system.

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

Atomic Transaction 2PC Protocol (Participant View)

Inbound Events	States						
	None	Active	Preparing	Prepared	PreparedSuccess	Committing	Aborting
Register Response	<i>Register Subordinate</i> Active	<i>Ignore</i> Active	<i>Ignore</i> Preparing	<i>Ignore</i> Prepared	<i>Ignore</i> PreparedSuccess	<i>Ignore</i> Committing	<i>Ignore</i> Aborting
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	<i>Resend Aborted, and Forget</i> Aborting
Commit	<i>Send Committed</i> None	<i>Invalid State</i> Aborting	<i>Invalid State</i> Aborting	<i>Invalid State</i> Aborting	<i>Initiate Commit Decision</i> Committing	<i>Ignore</i> Committing	<i>InconsistentInternalState</i> Aborting
Rollback	<i>Send Aborted</i> None	<i>Initiate Rollback, Send Aborted, and Forget</i> Aborting	<i>Initiate Rollback, Send Aborted, and Forget</i> Aborting	<i>Initiate Rollback, Send Aborted, and Forget</i> Aborting	<i>Initiate Rollback, Send Aborted, and Forget</i> Aborting	<i>InconsistentInternalState</i> Committing	<i>Send Aborted, and Forget</i> Aborting
Internal Events							
Expires Times Out	<i>N/A</i>	<i>Send Aborted</i> Aborting	<i>Send Aborted</i> Aborting	<i>Ignore</i> Prepared	<i>Ignore</i> PreparedSuccess	<i>Ignore</i> Committing	<i>Ignore</i> Aborting
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>	<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 and Forget</i> Committing	<i>N/A</i>
Rollback Decision	<i>N/A</i>	<i>N/A</i>	<i>Send Aborted</i> Aborting	<i>N/A</i>	<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>	<i>N/A</i>
Write Failed	<i>N/A</i>	<i>N/A</i>	<i>N/A</i>	<i>Initiate Rollback, Send Aborted, and Forget</i> Aborting	<i>N/A</i>	<i>N/A</i>	<i>N/A</i>
All Forgotten	None	<i>N/A</i>	<i>Send ReadOnly</i> None	<i>N/A</i>	<i>N/A</i>	None	None

- 511
- 512 Notes:
- 513 1. Transitions with a "N/A" as their action are inexpressible. A TM should view these transitions as
- 514 serious internal consistency issues, and probably fatal.
- 515 2. Internal events are those that are created either within a TM itself, or on its local system.

516

Appendix A. Acknowledgements

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

522

523 The following individuals have provided invaluable input into the initial contribution: Francisco Curbera,
524 IBM, Doug Davis, IBM, Gert Drapers, Microsoft, Don Ferguson, IBM, Kirill Gavrylyuk, Microsoft, Dan
525 House, IBM, Oisin Hurley, IONA, Thomas Mikalsen, IBM, Jagan Peri, Microsoft, John Shewchuk,
526 Microsoft, Stefan Tai, IBM.

527

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

530 **Participants:**

531 [Participant Name, Affiliation | Individual Member]

532 [Participant Name, Affiliation | Individual Member]

533

Appendix B. Revision History

Revision	yy-mm-dd	Editor	Changes Made
01	05-11-22	Mark Little	Initial Working Draft
01	05-11-23	Andrew Wilkinson	Initial Working Draft
02	06-02-12	Mark Little	Updated for issue i017
03	06-03-02	Andrew Wilkinson	Updated for issue 015
04	06-03-10	Andrew Wilkinson	Updated for issue 009
cd-01	06-03-15	Andrew Wilkinson	Updates to produce CD-01

534