



Web Services Atomic Transaction (WS-AtomicTransaction) Version 1.1

OASIS Standard incorporating Approved Errata

12 July 2007

Specification URIs:

This Version:

<http://docs.oasis-open.org/ws-tx/wstx-wsat-1.1-spec-errata-os/wstx-wsat-1.1-spec-errata-os.html>
<http://docs.oasis-open.org/ws-tx/wstx-wsat-1.1-spec-errata-os.doc>
<http://docs.oasis-open.org/ws-tx/wstx-wsat-1.1-spec-errata-os.pdf>

Previous Version:

<http://docs.oasis-open.org/ws-tx/wstx-wsat-1.1-spec-os/wstx-wsat-1.1-spec-os.html>
<http://docs.oasis-open.org/ws-tx/wstx-wsat-1.1-spec-os.doc>
<http://docs.oasis-open.org/ws-tx/wstx-wsat-1.1-spec-os.pdf>

Latest Approved Version:

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

Technical Committee:

OASIS Web Services Transaction (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>

Declared XML Namespaces:

<http://docs.oasis-open.org/ws-tx/wsat/2006/06>

Abstract:

The WS-AtomicTransaction specification provides the definition of the Atomic Transaction coordination type that is to be used with the extensible coordination framework described in WS-Coordination. This 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 was last revised or approved by the WS-TX TC on the above date. The level of approval is also listed above. Check the "Latest Approved Version" location noted above for possible later revisions of this document.

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

Copyright © OASIS Open 2007. All Rights Reserved.

All capitalized terms in the following text have the meanings assigned to them in the OASIS Intellectual Property Rights Policy (the "OASIS IPR Policy"). The full Policy may be found at the OASIS website.

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 section are included on all such copies and derivative works. However, this document itself may not be modified in any way, including by removing the copyright notice or references to OASIS, except as needed for the purpose of developing any document or deliverable produced by an OASIS Technical Committee (in which case the rules applicable to copyrights, as set forth in the OASIS IPR Policy, 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 OWNERSHIP RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

OASIS requests that any OASIS Party or any other party that believes it has patent claims that would necessarily be infringed by implementations of this OASIS Committee Specification or OASIS Standard, to notify OASIS TC Administrator and provide an indication of its willingness to grant patent licenses to such patent claims in a manner consistent with the IPR Mode of the OASIS Technical Committee that produced this specification.

OASIS invites any party to contact the OASIS TC Administrator if it is aware of a claim of ownership of any patent claims that would necessarily be infringed by implementations of this specification by a patent holder that is not willing to provide a license to such patent claims in a manner consistent with the IPR Mode of the OASIS Technical Committee that produced this specification. OASIS may include such claims on its website, but disclaims any obligation to do so.

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' procedures with respect to rights in any document or deliverable produced by an OASIS Technical Committee can be found on 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 OASIS Committee Specification or OASIS Standard, can be obtained from the OASIS TC Administrator. OASIS makes no representation that any information or list of intellectual property rights will at any time be complete, or that any claims in such list are, in fact, Essential Claims.

Table of contents

1	Introduction	5
1.1	Composable Architecture.....	5
1.2	Terminology	5
1.3	Namespace	6
1.3.1	Prefix Namespace	6
1.4	XSD and WSDL Files.....	6
1.5	Protocol Elements	6
1.6	Normative References	7
2	Atomic Transaction Context.....	9
3	Atomic Transaction Protocols	10
3.1	Preconditions	10
3.2	Completion Protocol.....	10
3.3	Two-Phase Commit Protocol	11
3.3.1	Volatile Two-Phase Commit Protocol	11
3.3.2	Durable Two-Phase Commit Protocol	12
3.3.3	2PC Diagram and Notifications.....	12
4	Policy Assertion	14
4.1	Assertion Model	14
4.2	Normative Outline	14
4.3	Assertion Attachment.....	14
4.4	Assertion Example	14
5	Transaction Faults	16
5.1	Inconsistent Internal State.....	17
5.2	Unknown Transaction	17
6	Security Model	18
7	Security Considerations.....	20
8	Use of WS-Addressing Headers.....	22
9	State Tables.....	23
9.1	Completion Protocol.....	23
9.2	2PC Protocol	24
A.	Acknowledgements.....	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.

WS-Coordination [[WSCOOR](#)] 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. WS-AtomicTransaction 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 WS-Coordination [[WSCOOR](#)] which defines the framework for the Atomic Transaction 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 by a transaction participant prior to commit are only tentative; typically they are neither persistent nor made visible outside the transaction. When an application finishes working on a transaction, 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 directs the participants to make the tentative actions final so they may, for example, be made persistent and be made visible outside the transaction. Abort directs the participants to make the tentative actions appear as if they 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.

This specification defines protocols that govern the outcome of Atomic Transactions. It is expected that existing transaction processing systems will use WS-AtomicTransaction to 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 [[WSCOOR](#)], WS-Security [[WSSec](#)]) 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 [[RFC2119](#)].

- 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
 - 54 XSD).
 - 55 • Characters are appended to elements, attributes, and <!-- descriptions --> as follows: "?" (0 or 1),
 - 56 "*" (0 or more), "+" (1 or more). The characters "[" and "]" are used to indicate that contained
 - 57 items are to be treated as a group with respect to the "?", "*", or "+" characters.
 - 58 • The XML namespace prefixes (defined below) are used to indicate the namespace of the element
 - 59 being defined.
 - 60 • Examples starting with <?xml contain enough information to conform to this specification; others
 - 61 examples are fragments and require additional information to be specified in order to conform.

62 1.3 Namespace

63 The XML namespace [XML-ns] URI that MUST be used by implementations of this specification is:

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

65 This MUST also be used as the CoordinationContext type for Atomic Transactions.

66 1.3.1 Prefix Namespace

67 The following namespaces are used in this document:

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/wsat/2006/06
wsa	http://www.w3.org/2005/08/addressing

68 1.4 XSD and WSDL Files

69 Dereferencing the XML namespace defined in section 1.3 will produce the Resource Directory
70 Description Language (RDDL) [RDDL] document that describes this namespace, including the XML
71 schema [XML-Schema1] [XML-Schema2] and WSDL [WSDL] declarations associated with this
72 specification.

73 SOAP bindings for the WSDL [WSDL], referenced in the RDDL [RDDL] document, MUST use "document"
74 for the *style* attribute.

75 1.5 Protocol Elements

76 The protocol elements define various extensibility points that allow other child or attribute content.
77 Additional children and/or attributes MAY be added at the indicated extension points but MUST NOT

78 contradict the semantics of the parent and/or owner, respectively. If a receiver does not recognize an
79 extension, the receiver SHOULD ignore the extension.

80 1.6 Normative References

- 81 **[RDDL]** Jonathan Borden, Tim Bray, eds. "Resource Directory Description Language
82 (RDDL) 2.0", <http://www.openhealth.org/RDDL/20040118/rddl-20040118.html>,
83 January 2004
- 84 **[RFC2119]** S. Bradner, "Key words for use in RFCs to Indicate Requirement Levels",
85 <http://www.ietf.org/rfc/rfc2119.txt>, IETF RFC2119, March 1997
- 86 **[SOAP11]** W3C Note, "SOAP: Simple Object Access Protocol 1.1",
87 <http://www.w3.org/TR/2000/NOTE-SOAP-20000508>, 08 May 2000
- 88 **[SOAP12]** W3C Recommendation, "SOAP Version 1.2 Part 1: Messaging Framework",
89 <http://www.w3.org/TR/soap12-part1>, June 2003
- 90 **[WSADDR]** Web Services Addressing (WS-Addressing) 1.0,
91 <http://www.w3.org/2005/08/addressing>, W3C Recommendation, May 2006
- 92 **[WSCOOR]** Web Services Coordination (WS-Coordination) 1.1, [http://docs.oasis-](http://docs.oasis-open.org/ws-tx/wscoor/2006/06)
93 [open.org/ws-tx/wscoor/2006/06](http://docs.oasis-open.org/ws-tx/wscoor/2006/06), OASIS, March 2006
- 94 **[WSDL]** Web Services Description Language (WSDL) 1.1,
95 <http://www.w3.org/TR/2001/NOTE-wsdl-20010315>
- 96 **[WSPOLICY]** Web Services Policy 1.2 – Framework (WS-Policy),
97 <http://www.w3.org/Submission/2006/SUBM-WS-Policy-20060425/>, W3C
98 Member Submission, 25 April 2006.
- 99 **[WSPOLICYATTACH]** Web Services Policy 1.2 – Attachment (WS-PolicyAttachment),
100 <http://www.w3.org/Submission/2006/SUBM-WS-PolicyAttachment-20060425/>,
101 W3C Member Submission, 25 April 2006.
- 102 **[WSSec]** OASIS Standard 200401, "Web Services Security: SOAP Message Security 1.0
103 (WS-Security 2004)", [http://docs.oasis-open.org/wss/2004/01/oasis-200401-](http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-soap-message-security-1.0.pdf)
104 [wss-soap-message-security-1.0.pdf](http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-soap-message-security-1.0.pdf), March 2004
- 105 **[WSSecConv]** Web Services Secure Conversation Language (WS-SecureConversation),
106 <http://schemas.xmlsoap.org/ws/2005/02/sc>, OpenNetwork, Layer7, Netegrity,
107 Microsoft, Reactivity, IBM, VeriSign, BEA Systems, Oblix, RSA Security, Ping
108 Identity, Westbridge, Computer Associates, February 2005
- 109 **[WSSecPolicy]** Web Services Security Policy Language (WS-SecurityPolicy),
110 <http://schemas.xmlsoap.org/ws/2005/07/securitypolicy>, Microsoft, VeriSign,
111 IBM, RSA Security, July 2005
- 112 **[WSTrust]** Web Services Trust Language (WS-Trust), ,
113 <http://schemas.xmlsoap.org/ws/2005/02/trust>, OpenNetwork, Layer7, Netegrity,
114 Microsoft, Reactivity, VeriSign, IBM, BEA Systems, Oblix, RSA Security, Ping
115 Identity, Westbridge, Computer Associates, February 2005
- 116 **[XML]** W3C Recommendation, "Extensible Markup Language (XML) 1.0 (Fourth
117 Edition)", <http://www.w3.org/TR/2006/REC-xml-20060816>, 16 August 2006
- 118 **[XML-ns]** W3C Recommendation, "Namespaces in XML (Second Edition)",
119 <http://www.w3.org/TR/2006/REC-xml-names-20060816>, 16 August 2006
- 120 **[XML-Schema1]** W3C Recommendation, "XML Schema Part 1: Structures Second Edition",
121 <http://www.w3.org/TR/2004/REC-xmlschema-1-20041028>, 28 October 2004

122 **[XML-Schema2]**
123

W3C Recommendation, " XML Schema Part 2: Datatypes Second Edition",
<http://www.w3.org/TR/2004/REC-xmlschema-2-20041028>, 28 October 2004

124 2 Atomic Transaction Context

125 WS-AtomicTransaction builds on WS-Coordination [WSCOOR], which defines an Activation service, a
126 Registration service, and a CoordinationContext type. Example message flows and a complete
127 description of creating and registering for coordinated activities is found in WS-Coordination [WSCOOR].

128 The Atomic Transaction coordination context is a CoordinationContext type with the coordination type
129 defined in this section. Atomic Transaction application messages that propagate a coordination context
130 MUST use an Atomic Transaction coordination context. If these application messages use a SOAP
131 binding, the Atomic Transaction coordination context MUST flow as a SOAP header in the message.

132 WS-AtomicTransaction adds the following semantics to the CreateCoordinationContext operation on the
133 Activation service:

- 134 • If the request includes the CurrentContext element, the target coordinator is interposed as a
135 subordinate to the coordinator stipulated inside the CurrentContext element.
- 136 • If the request does not include a CurrentContext element, the target coordinator creates a new
137 transaction and acts as the root.

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

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

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

3 Atomic Transaction Protocols

145

146 This specification defines the following protocols for Atomic Transactions:

- 147 • **Completion:** The completion protocol initiates commit processing. Based on each protocol's
148 registered participants, the coordinator begins with Volatile 2PC and then proceeds through
149 Durable 2PC. The final result is signaled to the initiator.
- 150 • **Two-Phase Commit (2PC):** The 2PC protocol coordinates registered participants to reach a
151 commit or abort decision, and ensures that all participants are informed of the final result. The
152 2PC protocol has two variants:
 - 153 ○ **Volatile 2PC:** Participants managing volatile resources such as a cache register for
154 this protocol.
 - 155 ○ **Durable 2PC:** Participants managing durable resources such as a database register
156 for this protocol.

157 A participant MAY register for more than one of these protocols.

3.1 Preconditions

158

159 The correct operation of the protocols requires that a number of preconditions must be established prior
160 to the processing:

- 161 1. The source SHOULD have knowledge of the destination's policies, if any, and the source
162 SHOULD be capable of formulating messages that adhere to this policy.
- 163 2. If a secure exchange of messages is required, then the source and destination MUST have
164 appropriate security credentials (such as transport-level security credentials or security tokens) in
165 order to protect the messages.

3.2 Completion Protocol

166

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

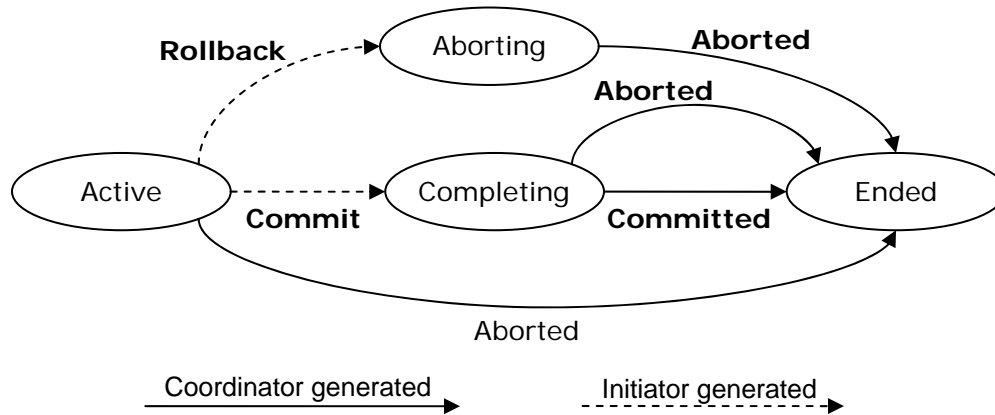
169 An initiator that registers for this protocol MUST use the following protocol identifier:

170

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

171 A Completion protocol coordinator MUST be the root coordinator of an Atomic Transaction. The
172 Registration service for a subordinate coordinator MUST respond to an attempt to register for this
173 coordination protocol with the WS-Coordination fault Cannot Register Participant.

174 The diagram below illustrates the protocol abstractly. Refer to section 9 State Tables for a detailed
175 description of this protocol.



176
 177 The coordinator accepts:
 178 Commit
 179 Upon receipt of this notification, the coordinator knows that the initiator has completed application
 180 processing. A coordinator that is Active SHOULD attempt to commit the transaction.

181 Rollback
 182 Upon receipt of this notification, the coordinator knows that the initiator has terminated application
 183 processing. A coordinator that is Active MUST abort the transaction.

184 The initiator accepts:
 185 Committed
 186 Upon receipt of this notification, the initiator knows that the coordinator reached a decision to
 187 commit.

188 Aborted
 189 Upon receipt of this notification, the initiator knows that the coordinator reached a decision to
 190 abort.

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

192 3.3 Two-Phase Commit Protocol

193 The Two-Phase Commit (2PC) protocol is a Coordination protocol that defines how multiple participants
 194 reach agreement on the outcome of an Atomic Transaction. The 2PC protocol has two variants: Volatile
 195 2PC and Durable 2PC.

196 3.3.1 Volatile Two-Phase Commit Protocol

197 Upon receiving a Commit notification in the Completion protocol, the root coordinator begins the prepare
 198 phase of all participants registered for the Volatile 2PC protocol. All participants registered for this
 199 protocol MUST respond before a Prepare is issued to a participant registered for Durable 2PC. Further
 200 participants MAY register with the coordinator until the coordinator issues a Prepare to any durable
 201 participant. Once this has happened the Registration Service for the coordinator MUST respond to any
 202 further Register requests with a Cannot Register Participant fault message. A volatile recipient is not
 203 guaranteed to receive a notification of the transaction's outcome.

204 Participants that register for this protocol MUST use the following protocol identifier:

205 `http://docs.oasis-open.org/ws-tx/wsac/2006/06/Volatile2PC`

206 **3.3.2 Durable Two-Phase Commit Protocol**

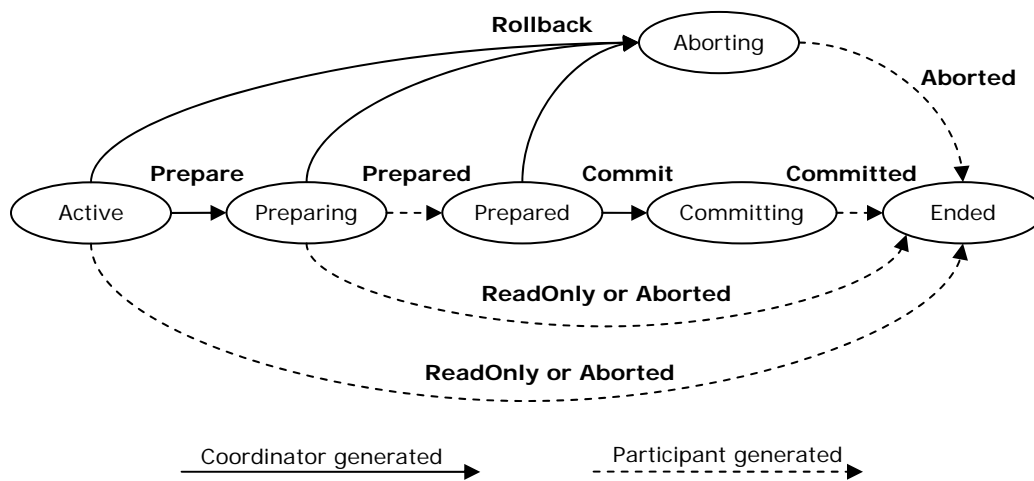
207 Upon successfully completing the prepare phase for Volatile 2PC participants, the root coordinator begins
208 the prepare phase for Durable 2PC participants. All participants registered for this protocol **MUST**
209 respond Prepared or ReadOnly before a Commit notification is issued to a participant registered for either
210 protocol.

211 Participants that register for this protocol **MUST** use the following protocol identifier:

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

213 **3.3.3 2PC Diagram and Notifications**

214 The diagram below illustrates the protocol abstractly. Refer to section 9 State Tables for a detailed
215 description of this protocol.



216
217 The participant accepts:

218 Prepare

219 Upon receipt of this notification, the participant knows to enter phase one and vote on the
220 outcome of the transaction. A participant that is Active **MUST** respond by sending Aborted,
221 Prepared, or ReadOnly notification as its vote. If the participant does not know of the transaction,
222 it **MUST** send an Aborted notification. If the participant knows that it has already voted, it **MUST**
223 resend the same vote.

224 Rollback

225 Upon receipt of this notification, the participant knows to abort and forget the transaction. A
226 participant that is not Committing **MUST** respond by sending an Aborted notification and
227 **SHOULD** then forget all knowledge of this transaction. If the participant does not know of the
228 transaction, it **MUST** send an Aborted notification to the coordinator.

229 Commit

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

233 The coordinator accepts:

234 Prepared

235 Upon receipt of this notification, the coordinator knows the participant is Prepared and votes to
236 commit the transaction.

237 ReadOnly

238 Upon receipt of this notification, the coordinator knows the participant votes to commit the
239 transaction, and has forgotten the transaction. The participant does not wish to participate in
240 phase two.

241 Aborted

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

244 Committed

245 Upon receipt of this notification, the coordinator knows the participant has committed and
246 forgotten the transaction.

247 Conforming implementations MUST implement the 2PC protocol.

248 4 Policy Assertion

249 WS-Policy Framework [WSPOLICY] and WS-Policy Attachment [WSPOLICYATTACH] collectively define
250 a framework, model and grammar for expressing the capabilities, requirements, and general
251 characteristics of entities in an XML Web services-based system. To enable a Web service to describe
252 transactional capabilities and requirements of a service and its operations, this specification defines an
253 Atomic Transaction policy assertion that leverages the WS-Policy [WSPOLICY] framework.

254 4.1 Assertion Model

255 The Atomic Transaction policy assertion is provided by a Web service to qualify the transactional
256 processing of messages associated with the particular operation to which the assertion is scoped. It
257 indicates whether a requester MAY or MUST include an Atomic Transaction coordination context flowed
258 with the message.

259 4.2 Normative Outline

260 The normative outline for the Atomic Transaction policy assertion is:

```
261 <wsat:ATAssertion [wsp:Optional="true"]? ... >  
262 ...  
263 </wsat:ATAssertion>
```

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

265 /wsat:ATAssertion

266 A policy assertion that specifies that an Atomic Transaction coordination context MUST be flowed
267 inside a requester's message. From the perspective of the requester, the target service that
268 processes the transaction MUST behave as if it had participated in the transaction. For application
269 messages that use a SOAP binding, the Atomic Transaction coordination context MUST flow as a
270 SOAP header in the message.

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

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

274 4.3 Assertion Attachment

275 Because the Atomic Transaction policy assertion indicates Atomic Transaction behavior for a single
276 operation, the assertion has an Operation Policy Subject [WSPOLICYATTACH].

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

- 279 • wsdl:portType/wsdl:operation – A policy expression containing the Atomic Transaction policy
280 assertion MUST NOT be attached to a wsdl:portType; the Atomic Transaction policy assertion
281 specifies a concrete behavior whereas the wsdl:portType is an abstract construct.
- 282 • wsdl:binding/wsdl:operation – A policy expression containing the Atomic Transaction policy
283 assertion SHOULD be attached to a wsdl:binding.

284 4.4 Assertion Example

285 An example use of the Atomic Transaction policy assertion follows:

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

```

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

```

309 Lines 8-11 are a policy expression that includes an Atomic Transaction policy assertion (line 9) to indicate
310 that an Atomic Transaction in WS-Coordination [WSCOOR] format MAY be used.

311 Lines 13-19 are a WSDL [WSDL] binding. Line 16 indicates that the policy in lines 8-11 applies to this
312 binding, specifically indicating that an Atomic Transaction MAY flow inside messages.

313 5 Transaction Faults

314 Atomic Transaction faults MUST include, as the [action] property, the following fault action URI:

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

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

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

320 [Code] The fault code.

321 [Subcode] The fault subcode.

322 [Reason] A human readable explanation of the fault.

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

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

326

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

327

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

```
329 <S12:Envelope>
330 <S12:Header>
331   <wsa:Action>
332     http://docs.oasis-open.org/ws-tx/wsat/2006/06/fault
333   </wsa:Action>
334   <!-- Headers elided for clarity. -->
335 </S12:Header>
336 <S12:Body>
337 <S12:Fault>
338   <S12:Code>
339     <S12:Value>[Code]</S12:Value>
340     <S12:Subcode>
341       <S12:Value>[Subcode]</S12:Value>
342     </S12:Subcode>
343   </S12:Code>
344   <S12:Reason>
345     <S12:Text xml:lang="en">[Reason]</S12:Text>
346   </S12:Reason>
347   <S12:Detail>
348     [Detail]
349     ...
350   </S12:Detail>
351 </S12:Fault>
352 </S12:Body>
353 </S12:Envelope>
```

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

```
355 <S11:Envelope>
356 <S11:Body>
357 <S11:Fault>
358   <faultcode> [Subcode] </faultcode>
```



```
359     <faultstring xml:lang="en">[Reason]</faultstring>
360     </S11:Fault>
361     </S11:Body>
362 </S11:Envelope>
```

363 5.1 Inconsistent Internal State

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

367 Properties:

368 **[Code]** Sender

369 **[Subcode]** wsat:InconsistentInternalState

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

371 **[Detail]** Unspecified

372 5.2 Unknown Transaction

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

375 Properties:

376 **[Code]** Sender

377 **[Subcode]** wsat:UnknownTransaction

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

379 **[Detail]** Unspecified

380

6 Security Model

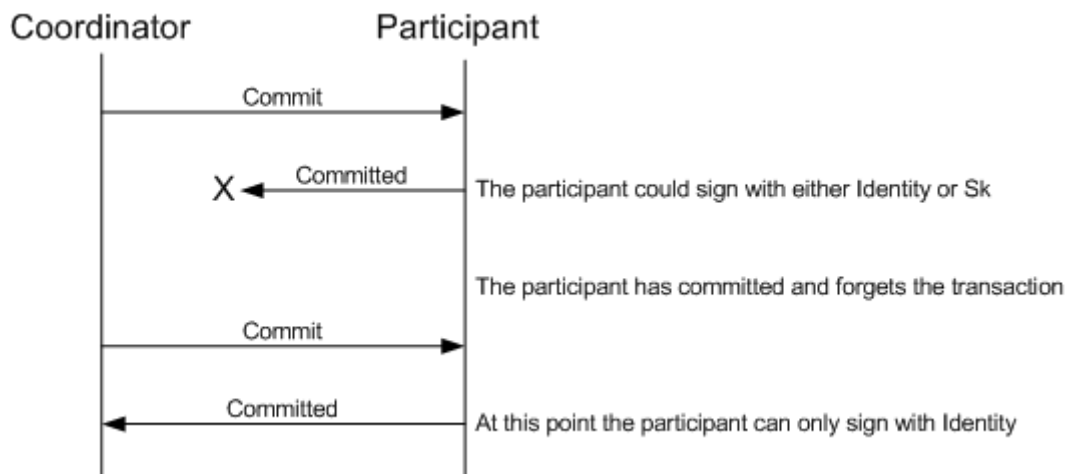
381 The security model for Atomic Transactions builds on the model defined in WS-Coordination [WSCOOR].
382 That is, services have policies specifying their requirements and requestors provide claims (either implicit
383 or explicit) and the requisite proof of those claims. Coordination context creation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

385 Because Atomic Transactions represent a specific use case rather than the general nature of
386 coordination contexts, additional aspects of the security model can be specified.

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

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

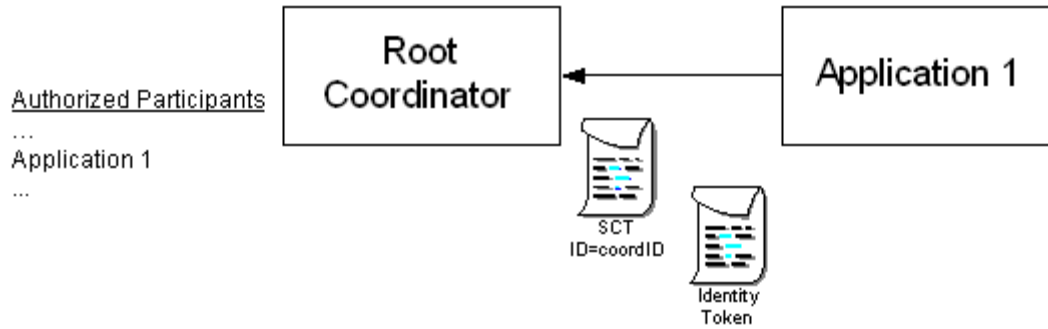
395 The participant can only prove its identity to the coordinator when it indicates that the specified
396 transaction is not in its log and assumed committed. This is illustrated in the figure below:



397

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

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



405

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

409

7 Security Considerations

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

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

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

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

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

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

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

- 436 • **Message alteration** – Alteration is prevented by including signatures of the message information
437 using WS-Security [WSSec].
- 438 • **Message disclosure** – Confidentiality is preserved by encrypting sensitive data using WS-
439 Security [WSSec].
- 440 • **Key integrity** – Key integrity is maintained by using the strongest algorithms possible (by
441 comparing secured policies – see WS-Policy [WSPOLICY] and WS-SecurityPolicy
442 [WSSecPolicy]).
- 443 • **Authentication** – Authentication is established using the mechanisms described in WS-Security
444 and WS-Trust [WSTrust]. Each message is authenticated using the mechanisms described in
445 WS-Security [WSSec].
- 446 • **Accountability** – Accountability is a function of the type of and string of the key and algorithms
447 being used. In many cases, a strong symmetric key provides sufficient accountability. However, in
448 some environments, strong PKI signatures are required.
- 449 • **Availability** – Many services are subject to a variety of availability attacks. Replay is a common
450 attack and it is RECOMMENDED that this be addressed as described in the next bullet. Other
451 attacks, such as network-level denial of service attacks are harder to avoid and are outside the
452 scope of this specification. That said, care should be taken to ensure that minimal processing be
453 performed prior to any authenticating sequences.

- 454
- 455
- 456
- 457
- **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.

458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493

8 Use of WS-Addressing Headers

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

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

The following statements define addressing interoperability requirements for the Atomic Transaction message types:

Non-terminal notification messages:

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

Both terminal and non-terminal notification messages:

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

Notification messages used in WS-AtomicTransaction protocols MUST include as the [action] property an action URI that consists of the wsat namespace URI concatenated with the "/" character and the element name of the message. For example:

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

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

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

All messages are delivered using connections initiated by the sender.

494 9 State Tables

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

497 Each cell in the tables uses the following convention:

498

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

499

500 Each state supports a number of possible events. Expected events are processed by taking the
 501 prescribed action and transitioning to the next state. Unexpected protocol messages MUST result in a
 502 fault message as defined in the state tables. These faults use standard fault codes as defined in either
 503 WS-Coordination [WSCOOR] or in section 5 Transaction Faults. Events that may not occur in a given
 504 state are labeled as N/A.

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 that are likely fatal conditions.
- 508 2. The "Internal events" shown are those events, created either within a TM itself or on its local
 509 system, that cause state changes and/or trigger the sending of a protocol message.

510 9.1 Completion Protocol

511

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

512

513 **9.2 2PC Protocol**

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

517

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	N/A	<i>Send Prepare</i> Preparing	N/A	N/A	N/A	N/A	N/A
User Rollback	N/A	<i>Send Rollback</i> Aborting	N/A	N/A	N/A	N/A	N/A
Expires Times Out	N/A	<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	N/A	N/A	<i>Resend Prepare</i> Preparing	N/A	N/A	<i>Resend Commit</i> Committing	N/A
Commit Decision	N/A	N/A	N/A	<i>Record Outcome</i> PreparedSuccess	N/A	N/A	N/A
Rollback Decision	N/A	<i>Send Rollback</i> Aborting	<i>Send Rollback</i> Aborting	<i>Send Rollback</i> Aborting	N/A	N/A	N/A
Write Done	N/A	N/A	N/A	N/A	<i>Send Commit</i> Committing	N/A	N/A
Write Failed	N/A	N/A	N/A	N/A	<i>Send Rollback</i> Aborting	N/A	N/A
Participant Abandoned	N/A	N/A	N/A	N/A	N/A	Durable: N/A Volatile: None	None

518

519 “Forget” implies that the subordinate’s participation is removed from the coordinator (if necessary), and
 520 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 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>

521

522

A. Acknowledgements

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

529

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

534

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

536 **Participants:**

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