



Web Services Atomic Transaction (WS-AtomicTransaction) 1.1

Committee Specification, 04 December 2006

Document Identifier:

wstx-wsat-1.1-spec-cs-01

Location:

<http://docs.oasis-open.org/ws-tx/wstx-wsat-1.1-spec-cs-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 was last revised or approved by the WS-TX TC on the above date. The level of approval is also listed above. Check the current location noted above for possible later revisions of this document. This document is updated periodically on no particular schedule.

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

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

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

Notices

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

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

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

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

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

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

Table of contents

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

1 Introduction

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

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

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

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

Atomic transactions have an all-or-nothing property. The actions taken 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 [KEYWORDS].

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

63 [[WSDL](#)].

64 1.3 Namespace

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

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

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

68 1.3.1 Prefix Namespace

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

69 1.4 XSD and WSDL Files

70 The XML schema and the WSDL declarations defined in this document can be found at the following

71 locations:

- 72 • <http://docs.oasis-open.org/ws-tx/wsat/2006/06/wsat.xsd>
- 73 • <http://docs.oasis-open.org/ws-tx/wsat/2006/06/wsat.wsdl>

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

75 attribute.

76 1.5 AT Protocol Elements

77 The protocol elements define various extensibility points that allow other child or attribute content.

78 Additional children and/or attributes MAY be added at the indicated extension points but MUST NOT

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

81 **1.6 Normative References**

82 **[KEYWORDS]**

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

85 **[SOAP11]**

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

88 **[SOAP12]**

89 W3C Recommendation, "SOAP Version 1.2 Part 1: Messaging Framework",
90 <http://www.w3.org/TR/soap12-part1>, June 2003

91 **[WSADDR]**

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

94 **[WSCOOR]**

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

97 **[WSDL]**

98 Web Services Description Language (WSDL) 1.1, [http://www.w3.org/TR/2001/NOTE-wsdl-](http://www.w3.org/TR/2001/NOTE-wsdl-20010315)
99 [20010315](http://www.w3.org/TR/2001/NOTE-wsdl-20010315)

100 **[WSPOLICY]**

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

103 **[WSPOLICYATTACH]**

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

107 **[WSSec]**

108 OASIS Standard 200401, "Web Services Security: SOAP Message Security 1.0 (WS-Security
109 2004)", [http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-soap-message-security-](http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-soap-message-security-1.0.pdf)
110 [1.0.pdf](http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-soap-message-security-1.0.pdf), March 2004

111 **[WSSecConv]**

112 Web Services Secure Conversation Language (WS-SecureConversation),
113 <http://schemas.xmlsoap.org/ws/2005/02/sc>, OpenNetwork, Layer7, Netegrity, Microsoft,
114 Reactivity, IBM, VeriSign, BEA Systems, Oblix, RSA Security, Ping Identity, Westbridge,
115 Computer Associates, February 2005

116 **[WSSecPolicy]**

117 Web Services Security Policy Language (WS-SecurityPolicy),
118 <http://schemas.xmlsoap.org/ws/2005/07/securitypolicy>, Microsoft, VeriSign, IBM, RSA Security,
119 July 2005

120 **[WSTrust]**

121 Web Services Trust Language (WS-Trust), , <http://schemas.xmlsoap.org/ws/2005/02/trust>,
122 OpenNetwork, Layer7, Netegrity, Microsoft, Reactivity, VeriSign, IBM, BEA Systems, Oblix, RSA
123 Security, Ping Identity, Westbridge, Computer Associates, February 2005

124 **[XML]**

125 W3C Recommendation, "Extensible Markup Language (XML) 1.0 (Fourth Edition)",
126 <http://www.w3.org/TR/2006/REC-xml-20060816>, 16 August 2006

127 **[XML-ns]**

128 W3C Recommendation, "Namespaces in XML (Second Edition)",
129 <http://www.w3.org/TR/2006/REC-xml-names-20060816>, 16 August 2006

130 **[XML-Schema1]**

131 W3C Recommendation, " XML Schema Part 1: Structures Second Edition",
132 <http://www.w3.org/TR/2004/REC-xmlschema-1-20041028>, 28 October 2004

133 **[XML-Schema2]**

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

136 2 Atomic Transaction Context

137 Atomic Transaction builds on WS-Coordination [WSCOOR], which defines an Activation service, a
138 Registration service, and a CoordinationContext type. Example message flows and a complete
139 description of creating and registering for coordinated activities is found in the WS-Coordination
140 specification [WSCOOR].

141 The Atomic Transaction coordination context is a CoordinationContext type with the coordination type
142 defined in this section. Application messages that propagate a transaction using the Atomic Transaction
143 protocol MUST use an Atomic Transaction coordination context. If these application messages use a
144 SOAP binding, the Atomic Transaction coordination context MUST flow as a SOAP header in the
145 message.

146 Atomic Transaction adds the following semantics to the CreateCoordinationContext operation on the
147 Activation service:

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

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

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

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

3 Atomic Transaction Protocols

159

160 This specification defines the following protocols for Atomic Transactions:

- 161 • **Completion:** The completion protocol initiates commit processing. Based on each protocol's
162 registered participants, the coordinator begins with Volatile 2PC and then proceeds through
163 Durable 2PC. The final result is signaled to the initiator.
- 164 • **Two-Phase Commit (2PC):** The 2PC protocol coordinates registered participants to reach a
165 commit or abort decision, and ensures that all participants are informed of the final result. The
166 2PC protocol has two variants:
 - 167 ○ **Volatile 2PC:** Participants managing volatile resources such as a cache register for
168 this protocol.
 - 169 ○ **Durable 2PC:** Participants managing durable resources such as a database register
170 for this protocol.

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

3.1 Preconditions

172

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

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

3.2 Completion Protocol

180

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

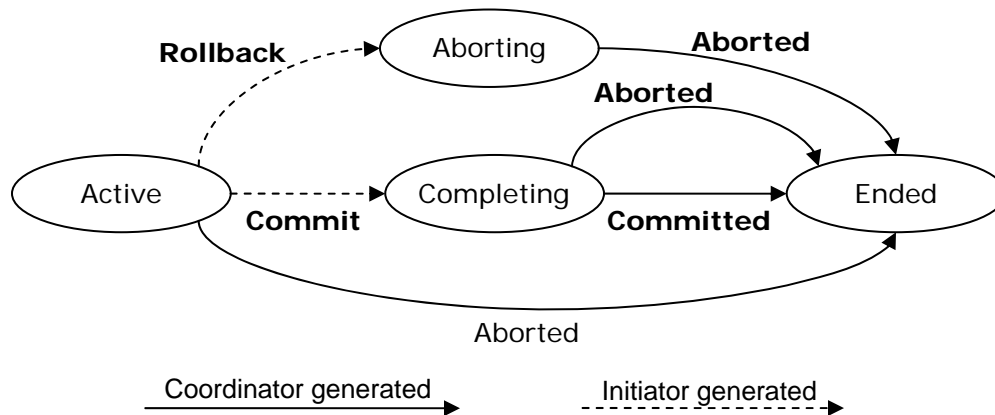
183 An initiator that registers for this protocol MUST use the following protocol identifier:

184

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

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

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



190
 191 The coordinator accepts:
 192 Commit
 193 Upon receipt of this notification, the coordinator knows that the initiator has completed application
 194 processing. A coordinator that is Active SHOULD attempt to commit the transaction.

195 Rollback
 196 Upon receipt of this notification, the coordinator knows that the initiator has terminated application
 197 processing. A coordinator that is Active MUST abort the transaction.

198 The initiator accepts:
 199 Committed
 200 Upon receipt of this notification, the initiator knows that the coordinator reached a decision to
 201 commit.

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

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

206 3.3 Two-Phase Commit Protocol

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

210 3.3.1 Volatile Two-Phase Commit Protocol

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

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

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

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

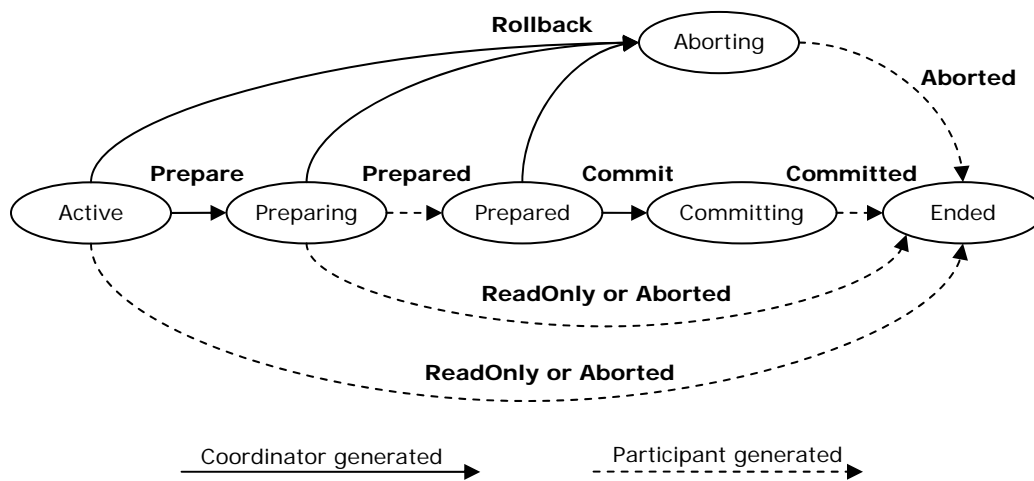
221 Upon successfully completing the prepare phase for Volatile 2PC participants, the root coordinator begins
222 the prepare phase for Durable 2PC participants. All participants registered for this protocol **MUST**
223 respond Prepared or ReadOnly before a Commit notification is issued to a participant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

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

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

227 **3.3.3 2PC Diagram and Notifications**

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



230
231 The participant accepts:

232 Prepare

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

238 Rollback

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

243 Commit

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

247 The coordinator accepts:

248 Prepared

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

251 ReadOnly

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

255 Aborted

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

258 Committed

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

261 Conforming implementations MUST implement the 2PC protocol.

262 4 AT Policy Assertion

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

268 4.1 Assertion Model

269 The AT policy assertion is provided by a web service to qualify the transactional processing of messages
270 associated with the particular operation to which the assertion is scoped. The AT policy assertion
271 indicates whether a requester MAY or MUST include an Atomic Transaction coordination context flowed
272 with the message.

273 4.2 Normative Outline

274 The normative outline for the AT policy assertion is:

```
275 <wsat:ATAssertion [wsp:Optional="true"]? ... >  
276 ...  
277 </wsat:ATAssertion>
```

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

279 /wsat:ATAssertion

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

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

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

288 4.3 Assertion Attachment

289 Because the AT policy assertion indicates Atomic Transaction behavior for a single operation, the
290 assertion has Operation Policy Subject [WSPOLICYATTACH].

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

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

297 4.4 Assertion Example

298 An example use of the AT policy assertion follows:

```
299 (01) <wsdl:definitions
```

```

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

```

321

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

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

326 5 Transaction Faults

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

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

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

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

333 [Code] The fault code.

334 [Subcode] The fault subcode.

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

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

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

339

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

340

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

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

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

```
368 <S11:Envelope>
369   <S11:Body>
370     <S11:Fault>
371       <faultcode> [Subcode] </faultcode>
```

```
372     <faultstring xml:lang="en">[Reason]</faultstring>
373     </S11:Fault>
374     </S11:Body>
375 </S11:Envelope>
```

376 **5.1 Inconsistent Internal State**

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

380 Properties:

381 **[Code]** Sender

382 **[Subcode]** wsat:InconsistentInternalState

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

384 **[Detail]** Unspecified

385 **5.2 Unknown Transaction**

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

388 Properties:

389 **[Code]** Sender

390 **[Subcode]** wsat:UnknownTransaction

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

392 **[Detail]** Unspecified

393

6 Security Model

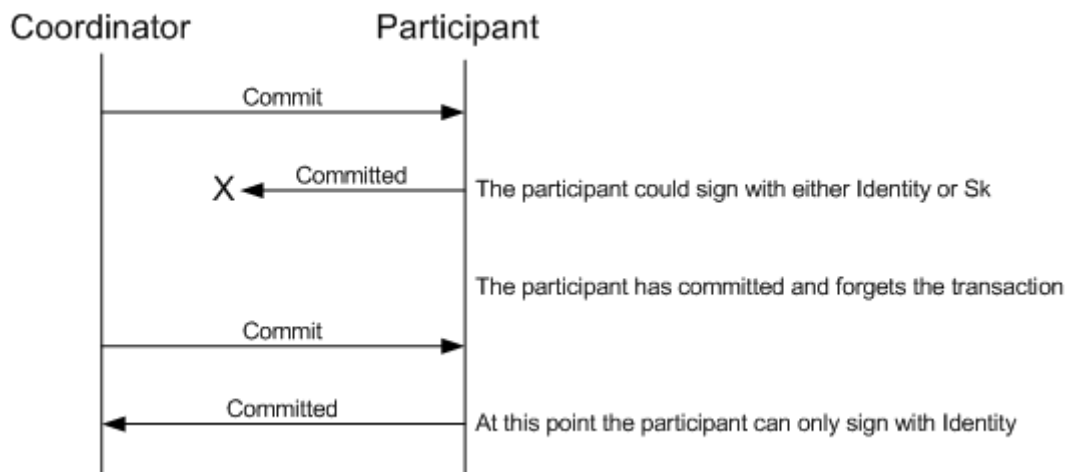
394 The security model for Atomic Transactions builds on the model defined in WS-Coordination [WSCOOR].
395 That is, services have policies specifying their requirements and requestors provide claims (either implicit
396 or explicit) and the requisite proof of those claims. Coordination context creation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

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

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

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


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



410

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

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



418

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

422

7 Security Considerations

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

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

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

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

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

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

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

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

- 467
- 468
- 469
- 470
- **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.

8 Use of WS-Addressing Headers

471
472 The protocols defined in WS-AtomicTransaction use a "one way" message exchange pattern consisting of
473 a sequence of notification messages between a Coordinator and a Participant. There are two types of
474 notification messages used in these protocols:

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

482 The following statements define addressing interoperability requirements for the WS-AtomicTransaction
483 message types:

484 Non-terminal notification messages:

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

488 Both terminal and non-terminal notification messages:

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

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

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

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

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

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

507 **9 State Tables**

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

510 Each cell in the tables uses the following convention:

511

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

512

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

518 Notes:

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

523 **9.1 Completion Protocol**

524

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

525

526 **9.2 2PC Protocol**

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

530

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

531
 532 “Forget” implies that the subordinate’s participation is removed from the coordinator (if necessary), and
 533 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>

534

535

A. Acknowledgements

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

542

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

547

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

549 **Participants:**

550 Martin Chapman, Oracle
551 Kevin Conner, JBoss Inc.
552 Paul Cotton, Microsoft Corporation
553 Doug Davis, IBM
554 Colleen Evans, Microsoft Corporation
555 Max Feingold, Microsoft Corporation
556 Thomas Freund, IBM
557 Robert Freund, Hitachi, Ltd.
558 Peter Furniss, Choreology Ltd.
559 Marc Goodner, Microsoft Corporation
560 Alastair Green, Choreology Ltd.
561 Daniel House, IBM
562 Ram Jeyaraman, Microsoft Corporation
563 Paul Knight, Nortel Networks Limited
564 Mark Little, JBoss Inc.
565 Jonathan Marsh, Microsoft Corporation
566 Monica Martin, Sun Microsystems
567 Joseph Fialli, Sun Microsystems
568 Eric Newcomer, IONA Technologies
569 Eisaku Nishiyama, Hitachi, Ltd.
570 Alain Regnier, Ricoh Company, Ltd.
571 Ian Robinson, IBM
572 Tom Rutt, Fujitsu Limited
573 Andrew Wilkinson, IBM