



Web Services ACID Specification (WS-ACID)

Editors draft version 0.2

Version created 5 July 2005

Editors

Mark Little (mark.little@arjuna.com)

Eric Newcomer (eric.newcomer@iona.com)

Greg Pavlik (greg.pavlik@oracle.com)

Copyright © 2005 The Organization for the Advancement of Structured Information
Standards [Appendix B]

41 An increasing number of applications are being constructed by combining or coordinating the
42 execution of multiple Web services, each of which may represent an interface to a different
43 underlying technology. The resulting applications can be very complex in structure, with complex
44 relationships between their constituent services. Furthermore, the execution of such an
45 application may take a long time to complete, and may contain long periods of inactivity, often
46 due to the constituent services requiring user interactions. In the loosely coupled environment
47 represented by Web services, long running applications will require support for recovery and
48 compensation, because machines may fail, processes may be cancelled, or services may be
49 moved or withdrawn. Web services transactions also must span multiple transaction models and
50 protocols native to the underlying technologies onto which the Web services are mapped.

51 A common technique for fault-tolerance is through the use of atomic transactions, which have the
52 well know ACID properties, operating on persistent (long-lived) objects. Transactions ensure that
53 only consistent state changes take place despite concurrent access and failures. However,
54 traditional transactions depend upon tightly coupled protocols, and thus are often not well suited
55 to more loosely-coupled Web services based applications, although they are likely to be used in
56 some of the constituent technologies. It is more likely that traditional transactions are used in the
57 minority of cases in which the cooperating Web services can take advantage of them, while new
58 mechanisms, such as compensation, replay, and persisting business process state, more suited
59 to Web services are developed and used for the more typical case.

60 WS-TXM provides a suite of transaction models, each suited to solving a different problem
61 domain. However, because WS-TXM leverages WS-CF, it is intended to allow flexibility in the
62 types of models supported. Therefore, if new models are required for other problem areas, they
63 can be incorporated within this specification.

64

Table of contents

66	1	Note on terminology	4
67	1.1	Namespace	4
68	1.1.1	Prefix Namespace	4
69	1.2	Referencing Specifications	4
70	2	Introduction	Error! Bookmark not defined.
71	2.1	Problem statement	Error! Bookmark not defined.
72	3	Architecture	Error! Bookmark not defined.
73	3.1	Invocation of Service Operations	5
74	3.2	Relationship to WSDL	6
75	3.3	Referencing and addressing conventions	6
76	4	WS-ACID	8
77	4.1	Interposition	Error! Bookmark not defined.
78	4.2	Restrictions imposed on using WS-CF	8
79	4.3	Two-phase commit	8
80	4.3.1	Coordinator state transitions for two-phase commit protocol	9
81	4.3.2	Two-phase participant state transitions	Error! Bookmark not defined.
82	4.3.3	Two-phase commit message interactions	10
83	4.3.4	Pre- and post- two-phase commit processing	13
84	4.3.5	Coordinator state transitions for synchronization protocol	14
85	4.3.6	Recovery and interposition	15
86	4.3.7	The context	15
87	4.3.8	Statuses	15
88	5	References	17

Deleted: 5

Deleted: 5

Deleted: 5

Deleted: 8

Deleted: 10

90 1 Note on terminology

91 The keywords "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD",
92 "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be
93 interpreted as described in RFC2119 [2].

94 Namespace URIs of the general form <http://example.org> and <http://example.com> represents some
95 application-dependent or context-dependent URI as defined in RFC 2396 [3].

96 1.1 Namespace

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

98 `http://docs.oasis-open.org/wscaf/2005/03/wsacid`

99 1.1.1 Prefix Namespace

Prefix	Namespace
wscf	http://docs.oasis-open.org/wscaf/2005/02/wscf
wsctx	http://docs.oasis-open.org/wscaf/2004/09/wsctx
wsacid	http://docs.oasis-open.org/wscaf/2005/07/wsacid
ref	http://docs.oasisopen.org/wsrn/2004/06/reference-1.1
wsdl	http://schemas.xmlsoap.org/wsdl/
xsd	http://www.w3.org/2001/XMLSchema
wsu	http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-wssecurity-utility-1.0.xsd
tns	targetNamespace

Comment: Kevin, can you check these are right (dates)?

100 1.2 Referencing Specifications

101 One or more other specifications may reference the WS-ACID specification. The usage of
102 optional items in WS-ACID is typically determined by the requirements of such as referencing
103 specification.

104 A referencing specification generally defines the protocol types based on WS-ACID. Any
105 application that uses WS-ACID must also decide what optional features are required. For the
106 purpose of this document, the term *referencing specification* covers both formal specifications
107 and more general applications that use WS-ACID.

108

109

2 Architecture

110

Atomic transactions are a well-known technique for guaranteeing consistency in the presence of failures [10]. The ACID properties of atomic transactions (Atomicity, Consistency, Isolation, and Durability) ensure that even in complex business applications consistency of state is preserved, despite concurrent accesses and failures. This is an extremely useful fault-tolerance technique, especially when multiple, possibly remote, resources are involved.

115

WS-ACID leverages the WS-CF and WS-Context specifications. Figure 4 illustrates the layering of WS-ACID onto WS-CF. WS-ACID defines a pluggable transaction protocol that can be used with the coordinator to negotiate a set of actions for all participants to execute based on the outcome of a series of related Web services executions. The executions are related through the use of shared context. Examples of coordinated outcomes include the classic two-phase commit protocol, a three phase commit protocol, open nested transaction protocol, asynchronous messaging protocol, or business process automation protocol.

116

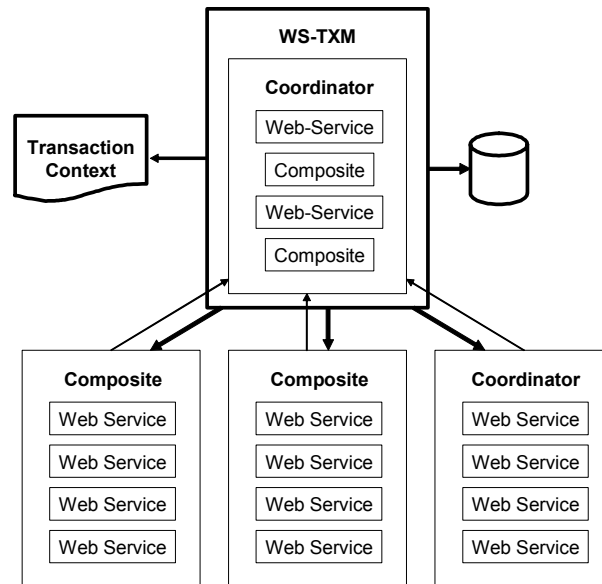
117

118

119

120

121



122

123

Figure 1, Relationship of transactions to coordination framework.

124

Coordinators can be participants of other coordinators, as shown above. When a coordinator registers itself with another coordinator, it can represent a series of local activities and map a neutral transaction protocol onto a platform-specific transaction protocol.

126

127

2.1 Invocation of Service Operations

128

How application services are invoked is outside the scope of this specification: they MAY use synchronous or asynchronous message passing.

129

130

Irrespective of how remote invocations occur, context information related to the sender's activity needs to be referenced or propagated. This specification determines the format of the context, how it is referenced, and how a context may be created.

131

132

133 In order to support both synchronous and asynchronous interactions, the components are
134 described in terms of the behavior and the interactions that occur between them. All interactions
135 are described in terms of correlated messages, which a referencing specification MAY abstract at
136 a higher level into request/response pairs.

137 Faults and errors that may occur when a service is invoked are communicated back to other Web
138 services in the activity via SOAP messages that are part of the standard protocol. To achieve this,
139 the fault mechanism of the underlying SOAP-based transport is used. For example, if an
140 operation fails because no activity is present when one is required, then the callback interface will
141 receive a SOAP fault including type of the fault and additional implementation specific information
142 items supported the SOAP fault definition. WS-Context specific fault types are described for each
143 operation. A fault type is communicated as an XML QName; the prefix consists of the WS-
144 Context namespace and the local part is the fault name listed in the operation description.

145 Note, a transientFault message is produced when the implementation finds it
146 cannot successfully execute the requested operation at that time from some
147 *temporary* reason. This reason may be implementation or referencing
148 specification specific. A receiver of a transientFault is free to retry the operation
149 which originally generated it on the assumption that eventually a different
150 response will be produced. Sub-types of transientFault MAY be further defined
151 using the fault model described which can allow for the communication of more
152 specific information on the type of fault.

153 As long as implementations ensure that the on-the-wire message formats are compliant with
154 those defined in this specification, how the end-points are implemented and how they expose the
155 various operations (e.g., via WSDL [1]) is not mandated by this specification. However, a
156 normative WSDL binding is provided by default in this specification.

157 Note, this specification does not assume that a reliable message delivery
158 mechanism has to be used for message interactions. As such, it MAY be
159 implementation dependant as to what action is taken if a message is not
160 delivered or no response is received.

161 **2.2 Relationship to WSDL**

162 Where WSDL is used in this specification it uses one-way messages with callbacks. This is the
163 normative style. Other binding styles are possible (perhaps defined by referencing specifications),
164 although they may have different acknowledgment styles and delivery mechanisms. It is beyond
165 the scope of WS-ACID to define these styles.

166 Note, conformant implementations MUST support the normative WSDL defined
167 in the specification where those respective interfaces are required. WSDL for
168 optional components in the specification is REQUIRED only in the cases where
169 the respective components are supported.

170 For clarity WSDL is shown in an abbreviated form in the main body of the document: only
171 portTypes are illustrated; a default binding to SOAP 1.1-over-HTTP is also assumed as per [1].

172 **2.3 Referencing and addressing conventions**

173 There are multiple mechanisms for addressing messages and referencing Web services currently
174 proposed by the Web services community. This specification defers the rules for addressing
175 SOAP messages to existing specifications; the addressing information is assumed to be placed in
176 SOAP headers and respect the normative rules required by existing specifications.

177 However, the Coordination Framework message set requires an interoperable mechanism for
178 referencing Web Services. For example, context structures may reference the service that is used
179 to manage the content of the context. To support this requirement, WS-CAF has adopted an open
180 content model for service references as defined by the Web Services Reliable Messaging
181 Technical Committee [5]. The schema is defined in [6][7] and is shown in Figure 1.

```

182 <xsd:complexType name="ServiceRefType">
183   <xsd:sequence>
184     <xsd:any namespace="##other" processContents="lax"/>
185   </xsd:sequence>
186   <xsd:attribute name="reference-scheme" type="xsd:anyURI"
187     use="optional"/>
188 </xsd:complexType>

```

189 Figure 2, service-ref Element

190 The ServiceRefType is extended by elements of the context structure as shown in Figure 2.

```

191 <xsd:element name="context-manager" type="ref:ServiceRefType"/>

```

192 Figure 3, ServiceRefType example.

193 Within the **ServiceRefType**, the reference-scheme is the namespace URI for the referenced
194 addressing specification. For example, the value for WSRef defined in the WS-MessageDelivery
195 specification [4] would be <http://www.w3.org/2004/04/ws-messagedelivery>. The value for WSRef
196 defined in the WS-Addressing specification [8] would be
197 <http://schemas.xmlsoap.org/ws/2004/08/addressing>. The reference scheme is optional and need
198 only be used if the namespace URI of the QName of the Web service reference cannot be used
199 to unambiguously identify the addressing specification in which it is defined.

200 Messages sent to referenced services **MUST** use the addressing scheme defined by the
201 specification indicated by the value of the reference-scheme element if present. Otherwise, the
202 namespace URI associated with the Web service reference element **MUST** be used to determine
203 the required addressing scheme. A service that requires a service reference element **MUST** use
204 the mustUnderstand attribute for the SOAP header element within which it is enclosed and **MUST**
205 return a mustUnderstand SOAP fault if the reference element isn't present and understood.

206 Note, it is assumed that the addressing mechanism used by a given
207 implementation supports a reply-to or sender field on each received message so
208 that any required responses can be sent to a suitable response endpoint. This
209 specification requires such support and does not define how responses are
210 handled.

211 To preserve interoperability in deployments that contain multiple addressing schemes, there are
212 no restrictions on a system, beyond those of the composite services themselves. However, it is
213 **RECOMMENDED** where possible that composite applications confine themselves to the use of
214 single addressing and reference model.

215 Because the prescriptive interaction pattern used by WS-ACID is based on one-way messages
216 with callbacks, it is possible that an endpoint may receive an unsolicited or unexpected message.
217 The recipient is free to do whatever it wants with such messages.

218

3 WS-ACID

219 The *ACID transaction* model recognizes that Web Services are for interoperability as much as for
220 the Internet. As such, interoperability of existing transaction processing systems will be an
221 important part of Web Services Transaction Management: such systems already form the
222 backbone of enterprise level applications and will continue to do so for the Web Services
223 equivalent. Business-to-business activities will typically involve back-end transaction processing
224 systems either directly or indirectly and being able to tie together these environments will be the
225 key to the successful take-up of Web Services transactions.

226 Although ACID transactions may not be suitable for all Web Services, they are most definitely
227 suitable for some, and particularly high-value interactions such as those involved in finance. As a
228 result, the ACID transaction model has been designed with interoperability in mind. Within this
229 model it is assumed that all services (and associated participants) provide ACID semantics and
230 that any use of atomic transactions occurs in environments and situations where this is
231 appropriate: in a trusted domain, over short durations.

232 In the ACID model, each activity is bound to the scope of a transaction, such that the end of an
233 activity automatically triggers the termination (commit or rollback) of the associated transaction.

234 The coordinator-type URI for the ACID transaction model is
235 <http://www.webservicestransactions.org/wsd/wstxm/tx-acid/2003/03>

Comment: Update.

3.1 Restrictions imposed on using WS-CF

237 As a Referencing Specification, the WS-ACID transaction model imposes the following
238 restrictions on using WS-CF:

- 239 • It is illegal to attempt to remove a participant from a transaction at any time. When the
240 transaction terminates, participants are implicitly removed. As such, any attempt to call
241 *removeParticipant* will result in the *wrongState* error message being returned.

3.2 Two-phase commit

242 The ACID transaction model uses a traditional two-phase commit protocol [2] with the following
243 optimizations:

- 244 • *Presumed rollback*: the transaction coordinator need not record information about the
245 participants in stable storage until it decides to commit, i.e., until after the prepare phase
246 has completed successfully.
- 247 • *One-phase*: if the coordinator discovers that only a single participant is registered then it
248 SHOULD omit the prepare phase..
- 249 • *Read-only*: a participant that is responsible for a service that did not modify any
250 transactional data during the course of the transaction can indicate to the coordinator
251 during prepare that it is a *read-only participant* and the coordinator SHOULD omit it from
252 the second phase of the commit protocol.

253 Participants that have successfully passed the *prepare* phase are allowed to make autonomous
254 decisions as to whether they commit or rollback. A participant that makes such an autonomous
255 choice *must* record its decision in case it is eventually contacted to complete the original
256 transaction. If the coordinator eventually informs the participant of the fate of the transaction and
257 it is the same as the autonomous choice the participant made, then there is obviously no
258 problem: the participant simply got there before the coordinator did. However, if the decision is
259 contrary, then a non-atomic outcome has happened: a *heuristic outcome*, with a corresponding
260 *heuristic decision*.

261 The possible heuristic outcomes are:

- 263 • *Heuristic rollback*: the commit operation failed because some or all of the participants
264 unilaterally rolled back the transaction.
- 265 • *Heuristic commit*: an attempted rollback operation failed because all of the participants
266 unilaterally committed. This may happen if, for example, the coordinator was able to
267 successfully prepare the transaction but then decided to roll it back (e.g., it could not
268 update its log) but in the meanwhile the participants decided to commit.
- 269 • *Heuristic mixed*: some updates were committed while others were rolled back.
- 270 • *Heuristic hazard*: the disposition of some of the updates is unknown. For those which are
271 known, they have either all been committed or all rolled back.

272 3.2.1 State transitions and relationship to WS-Context

273 WS-ACID is a referencing specification for WS-CF and hence leverages the *activity group*
274 concept. When an application creates a new activity group (by sending a **wsctx:begin** message
275 to the relevant Context Service), an associated WS-ACID coordinator MAY be created in the
276 Active state, as shown in [Figure 4](#).

Deleted: Figure 45

277 Note, participants enlisted with a WS-ACID activity group progress through the
278 same state transitions.

279 The coordinator has the lifetime period associated with the activity: if the activity timeout elapses
280 before the activity has terminated, then the transaction will be terminated in the RolledBack state.

281 A transactional activity can be terminated via the wsctx:complete message in one of two ways:

- 282 • Committed: the transaction commits.
- 283 • Rollback: the transaction rolls back.

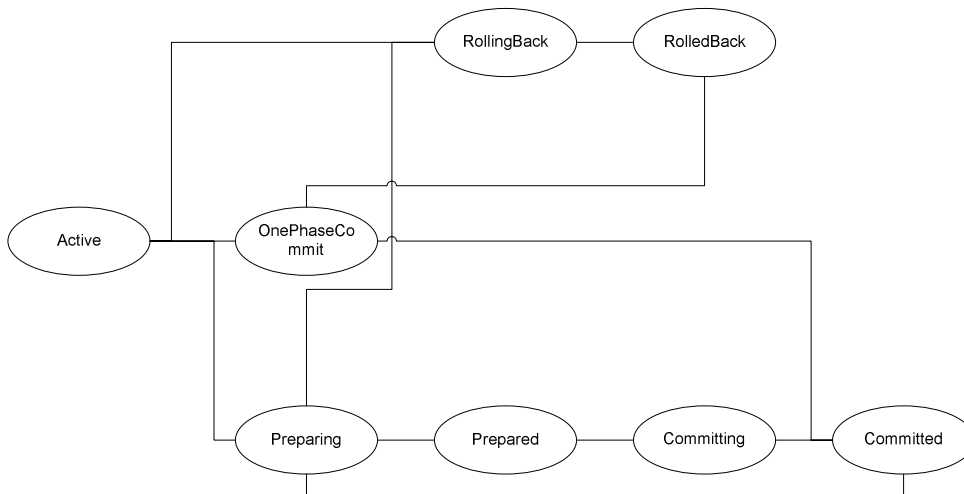
Comment: We need to decide what the commit/rollback data within the complete message looks like.

284 If the transaction is instructed to commit then the application sends an appropriate
285 **wsctx:complete** message to the Context Service. If there is only a single participant enrolled
286 with the transaction then the coordinator SHOULD use the one-phase commit optimization. As
287 such, the coordinator begins the *OnePhaseCommit* protocol and either transits to the *RolledBack*
288 or *Committed* state, depending upon the result returned by the participant. **The activity**
289 **completion status is either *Failure* or *Success* respectively.**

Comment: Open issue: how are errors communicated back to the application via the wsctx:completed message?

290 If there are multiple participants enrolled with the transaction, the coordinator transits to the
291 *Preparing* state and begins to execute the two-phase commit protocol by sending the
292 **wsacid:prepare** message to each participant. If all of the participants indicate that the services
293 they represent performed no work (i.e., are read only) then the transaction is complete and the
294 coordinator transits to the *Committed* state.

295 Any failures from a participant or indication that it cannot prepare cause the coordinator to
296 rollback (move to the *RollingBack* state) and send wsacid:rollback messages to all of the
297 participants. It then transits to the *RolledBack* state.



298

299

Figure 4, Transaction coordinator two-phase status transition.

300

Assuming all participants have prepared successfully, the transaction coordinator makes the decision as to whether to commit or rollback and must record sufficient information on stable storage to ensure this decision can be completed in the event of a failure. It is then in the *Prepared* state. When the coordinator starts the second phase of the commit protocol it is in the *Committing* state and ultimately moves to the *Committed* state.

301

302

303

304

305

3.2.2 Two-phase commit message interactions

306

In this section we shall describe the message exchanged between the coordinator and the participants. Although the text refers to the coordinator soliciting responses from participants, in some cases participants MAY send unsolicited responses to the coordinator; where this is the case it will be explicitly stated.

307

308

309

310

The ACID transaction model supports two styles of participant service implementation: the *singleton* approach, whereby one participant service (end-point) is implicitly associated with only one transaction, and the *factory* approach, whereby a single participant service may manage participants on behalf of many different transactions. Therefore, all operations on the participant service are associated with the current context, i.e., it is propagated to the participants in order to identify which transaction is to be operated on. The unique participant identification is also present on each message.

311

312

313

314

315

316

317

The two-phase commit sub-protocol URI is

318

<http://www.webservicestransactions.org/wsd/wstxm/tx-acid/2pc/2003/03> and this is used in the *addParticipant* message. An enlisted Participant Service should expect to receive the following messages (illustrated in [Figure 5](#)):

319

320

Comment: Need to update.

Deleted: Figure 58

321

- prepare: The coordinator is preparing. The participant can respond with a *voteReadOnly*, *voteCommit* or *voteRollback* messages indicating whether or not it is willing to commit. **If *voteCommit* is used then optional Qualifiers may be sent back to augment the protocol.** The *voteReadOnly* and *voteRollback* messages MAY be sent autonomously by the participant, i.e., before any **wsacid:prepare** message is received. However, the participant SHOULD be able to deal with a subsequent **wsacid:prepare** message. If an unreliable transport mechanism is used, then there may be an arbitrary number of these messages. If the participant is a subordinate coordinator and finds that it cannot determine the status of some of its enlisted participants then an error message with the **wsacid:HeuristicHazardFault** error code will be returned. Alternatively, if a subordinate

322

323

324

325

326

327

328

329

330

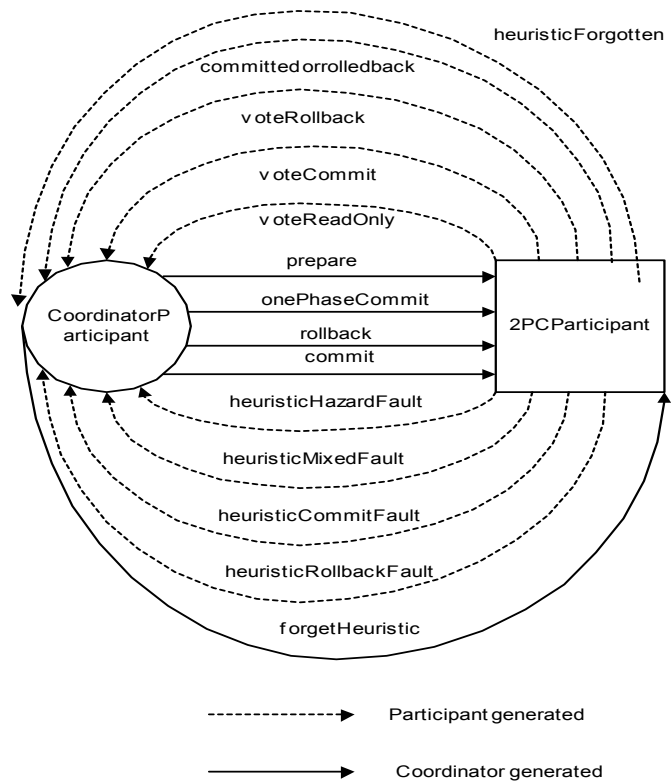
Comment: Issue 275

331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368

coordinator finds that some of the participants have committed and some have rolled back then it must return the **wsacid:HeuristicMixedFault** error message.

Comment: SOAP faults

- rollback: The coordinator is rolling back. If the participant is receiving this message after a **wsacid:prepare** message, then any error at this point will cause a heuristic outcome. If the participant is a subordinate coordinator and cannot determine how all of its enlisted participants terminated then it must return an error message with the **wsacid:HeuristicHazardFault** fault code. If the participant is a subordinate coordinator and some of its enlisted participants committed then it must return the **wsacid:HeuristicMixedFault** fault code. If the participant commits rather than rolls back then it must return the **wsacid:HeuristicCommitFault** message. Otherwise the participant sends the *rolledback* message. The **wsacid:rolledback** message MAY be sent autonomously by the participant, i.e., before any **wsacid:rollback** message is received. However, the participant SHOULD be able to deal with a subsequent **wsacid:rollback** message. If an unreliable transport mechanism is used, then there may be an arbitrary number of these messages.
- commit: The coordinator is top-level and is committing. Any error at this point will cause a heuristic outcome. If the participant is a subordinate coordinator and cannot determine how all of its enlisted participants terminated then it must return an error message with the **wsacid:HeuristicHazardFault** fault code. If the participant is a subordinate coordinator and some of its enlisted participants committed then it must return the **wsacid:HeuristicMixedFault** fault code. If the participant rolls back rather than commits then it must return the **wsacid:HeuristicRollbackFault** fault code. Otherwise the participant returns a *committed* message.
- onePhaseCommit: If only a single participant is registered with a two-phase coordinator then the coordinator SHOULD optimize the commit stage by not executing the prepare phase. If the participant is a subordinate coordinator and cannot determine how all of its enlisted participants terminated then it must return an error message with the **wsacid:HeuristicHazardFault** fault code. If the participant is a subordinate coordinator and some of its enlisted participants committed then it must return the **wsacid:HeuristicMixedFault** fault code. If the participant rolls back rather than commits then it must return the **wsacid:HeuristicRollbackFault** fault code. Otherwise the participant returns either the *committed* or *rolledback* message.
- forgetHeuristic: The participant made a post-prepare choice that was contrary to the coordinator's outcome. Hence it may have caused a non-atomic (heuristic) outcome. If this happens, the participant *must* remember the decision it took (persistently) until the coordinator tells it via this message that it is safe to forget. Success is indicated by sending the *heuristicForgotten* message. Any other response is assumed to indicate a failure.



369

370

Figure 5, AT coordinator-to-participant message exchanges.

371

The WSDL portType declarations for the CoordinatorParticipant and twoPCParticipant roles are shown in [Figure 6](#).

372

Comment: Need to change the names.
Deleted: Figure 69

373

```

<wsdl:portType name="twoPCParticipantPortType">
  <wsdl:operation name="prepare">
    <wsdl:input message="tns:PrepareMessage"/>
  </wsdl:operation>
  <wsdl:operation name="onePhaseCommit">
    <wsdl:input message="tns:OnePhaseCommitMessage"/>
  </wsdl:operation>
  <wsdl:operation name="rollback">
    <wsdl:input message="tns:RollbackMessage"/>
  </wsdl:operation>
  <wsdl:operation name="commit">
    <wsdl:input message="tns:CommitMessage"/>
  </wsdl:operation>
  <wsdl:operation name="forgetHeuristic">
    <wsdl:input message="tns:ForgetHeuristicMessage"/>
  </wsdl:operation>
</wsdl:portType>
<wsdl:portType name="CoordinatorParticipantPortType">
  <wsdl:operation name="committed">
    <wsdl:input message="tns:CommittedMessage"/>
  </wsdl:operation>
  <wsdl:operation name="rolledBack">
    <wsdl:input message="tns:RolledBackMessage"/>
  </wsdl:operation>
  <wsdl:operation name="vote">
  
```

374

375

376

377

378

379

380

381

382

383

384

385

386

387

388

389

390

391

392

393

394

395

396

397

```

398     <wsdl:input message="tns:VoteMessage"/>
399   </wsdl:operation>
400   <wsdl:operation name="heuristicForgotten">
401     <wsdl:input message="tns:HeuristicForgottenMessage"/>
402   </wsdl:operation>
403   <wsdl:operation name="heuristicFault">
404     <wsdl:input message="tns:HeuristicFaultMessage"/>
405   </wsdl:operation>
406 </wsdl:portType>

```

407 *Figure 6, WSDL portType Declarations for Coordinator and 2PCParticipant Roles*

408 Note, although an application Web Service may play the role of a participant, it is
 409 not required to.

410 3.3 Pre- and post- two-phase commit processing

411 Most modern transaction processing systems allow the creation of participants that do not take
 412 part in the two-phase commit protocol, but are informed before it begins and after it has
 413 completed. They are called Synchronizations, and are typically employed to flush volatile
 414 (cached) state, which may be being used to improve performance of an application, to a
 415 recoverable object or database prior to the transaction committing; once flushed, the data will be
 416 controlled by a two-phase aware participant.


417 The sub-protocol URI for the synchronization protocol is
 418 <http://www.webservicestransactions.org/wsd/wstxm/tx-acid/sync/2003/03> and this is used in the
 419 *addParticipant* invocation.

Comment: Needs updating.

420 The message exchanges (ignoring the normal WS-CF coordinator-to-participant message
 421 exchanges, including failures) are illustrated in [Figure 7](#):

Deleted: Figure 710

- 422 • *beforeCompletion*: A Synchronization participant is informed that the coordinator it is
 423 registered with is about to complete the two-phase protocol and in what state, i.e.,
 424 committing or rolling back. The failure of the participant at this stage will cause the
 425 coordinator to cancel if it is not already doing so.
- 426 • *afterCompletion*: A Synchronization participant is informed that the coordinator it is
 427 registered with has completed the two-phase protocol and in what state, i.e., committed
 428 or rolled back (via the associated Status). Any failures by the participant at this stage
 429 have no affect on the transaction.



430
 431 **Figure 7, AT coordinator-to-synchronization message exchanges.**

432 The WSDL portType declarations for the CoordinatorParticipant and Synchronization roles are
 433 shown in [Figure 8](#).

Deleted: Figure 811

```

434 <wsdl:portType name="SynchronizationPortType">
435   <wsdl:operation name="beforeCompletion">

```

```

436     <wsdl:input message="tns:BeforeCompletionMessage"/>
437   </wsdl:operation>
438   <wsdl:operation name="afterCompletion">
439     <wsdl:input message="tns:AfterCompletionMessage"/>
440   </wsdl:operation>
441 </wsdl:portType>
442 <wsdl:portType name="CoordinatorParticipantPortType">
443   <wsdl:operation name="beforeCompletionParticipantRegistered">
444     <wsdl:input
445 message="tns:BeforeCompletionParticipantRegisteredMessage"/>
446   </wsdl:operation>
447   <wsdl:operation name="afterCompletionParticipantRegistered">
448     <wsdl:input
449 message="tns:AfterCompletionParticipantRegisteredMessage"/>
450   </wsdl:operation>
451 </wsdl:portType>

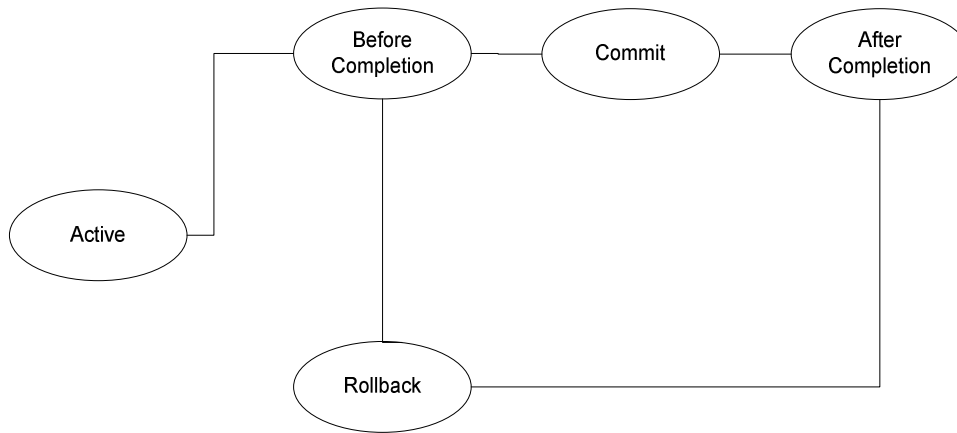
```

452 *Figure 8, WSDL portType Declarations for Coordinator and 2PCParticipant Roles.*

453 Note, the participant is registered for both beforeCompletion and
454 afterCompletion.

455 3.3.1 State transitions for synchronization protocol

456 The state transitions for the transaction coordinator which has enrolled Synchronizations is shown
457 in Figure 12. In this scenario we assume the transaction is committing: if it were to rollback, then
458 only the *AfterCompletion* message will be sent from the coordinator to the Synchronization
459 participants.



460
461 **Figure 9, Transaction coordinator Synchronization state transitions.**

462 The coordinator moves into the *BeforeCompletion* state and sends each enrolled Synchronization
463 the *beforeCompletion* message. Any error received by the coordinator from a Synchronization at
464 this stage will force the transaction to rollback. Assuming no errors occur, the two-phase commit
465 protocol is executed, as detailed previously. Once the protocol has completed, the coordinator

466 transits to the *AfterCompletion* status and sends the *afterCompletion* message to all
467 Synchronizations; any errors at this stage do not affect the transaction outcome and how they are
468 dealt with is implementation dependant.

469 3.4 Recovery and interposition

470 Because WS-ACID is a Referencing Specification of WS-CF, interposition is allowed though not
471 required. Individual participants may be subordinate coordinators to improve performance or to
472 federate a distributed environment into separate domains (possibly managed by different
473 organizations or transaction management systems).

474 Each participant or subordinate coordinator is responsible for ensuring that sufficient data is
475 made durable in order to complete the transaction in the event of failures. *Recovering participants
476 or coordinators use the recovery mechanisms defined in WS-CF to determine the current status
477 of a transaction/participant and act accordingly.* Interposition and check pointing of state allow the
478 system to drive a consistent view of the outcome and recovery actions taken, but allowing always
479 the possibility that recovery isn't possible and must be logged or flagged for the administrator.

480 Although enterprise transaction systems address the aspects of distributed recovery, in a large
481 scale environment or in the presence of long term failures, recovery may not be automatic. As
482 such, manual intervention may be necessary to restore an application's consistency.

483 3.5 The context

```
484 <xs:complexType name="ContextType">  
485   <xs:complexContent>  
486     <xs:extension base="wstxm:ContextType"/>  
487   </xs:complexContent>  
488 </xs:complexType>  
489 <xs:element name="context" type="tns:ContextType"/>
```

490 *Figure 10, Transaction Context.*

491 3.6 Statuses

492 The following extensions to the WS-Context Status type MAY be returned by participants and the
493 Context Service to indicate the outcome of executing relevant parts of the protocol and are also
494 used to indicate the current status of the transaction:

- 495 • RollbackOnly: the status of the coordinator or participant is that it will rollback eventually.
- 496 • RollingBack: the coordinator or participant is in the process of rolling back.
- 497 • RolledBack: the coordinator/participant has rolled back. This may be a transient and in
498 fact, because the protocol uses a presumed-abort optimisation, the NoActivity status can
499 be used to infer that the coordinator cancelled.
- 500 • Committing: the coordinator/participant is in the process of committing. This does not
501 mean that the final outcome will be Committed.
- 502 • Committed: the coordinator/participant has confirmed.
- 503 • HeuristicRollback: all of the participants rolled back when they were asked to commit.
- 504 • HeuristicCommit: all of the participants committed when they were asked to rollback.
- 505 • HeuristicHazard: some of the participants rolled back, some committed and the outcome
506 of others is indeterminate.
- 507 • HeuristicMixed: some of the participants rolled back whereas the remainder committed.
- 508 • Preparing: the coordinator/participant is preparing.
- 509 • Prepared: the coordinator/participant has prepared.

Comment: Issue – need to add a getStatus to the Participant Service WSDL?

510 | These are specified in the schema, as per [Figure 11](#),

Deleted: Figure 1114

```
511 <xs:simpleType name="StatusType">
512   <xs:restriction base="wstxm:StatusType">
513     <xs:enumeration value="activity.status.tx-acid.ROLLBACK_ONLY"/>
514     <xs:enumeration value="activity.status.tx-acid.ROLLING_BACK"/>
515     <xs:enumeration value="activity.status.tx-acid.ROLLED_BACK"/>
516     <xs:enumeration value="activity.status.tx-acid.COMMITTING"/>
517     <xs:enumeration value="activity.status.tx-acid.COMMITTED"/>
518     <xs:enumeration value="activity.status.tx-
519 acid.HEURISTIC_ROLLBACK"/>
520     <xs:enumeration value="activity.status.tx-acid.HEURISTIC_HAZARD"/>
521     <xs:enumeration value="activity.status.tx-acid.HEURISTIC_MIXED"/>
522     <xs:enumeration value="activity.status.tx-acid.PREPARING"/>
523     <xs:enumeration value="activity.status.tx-acid.PREPARED"/>
524   </xs:restriction>
525 </xs:simpleType>
```

526 | *Figure 11, StatusType.*

527

4 References

- 528 [1] WSDL 1.1 Specification, see <http://www.w3.org/TR/wsdl>
- 529 [2] "Key words for use in RFCs to Indicate Requirement Levels," RFC 2119, S. Bradner, Harvard
530 University, March 1997.
- 531 [3] "Uniform Resource Identifiers (URI): Generic Syntax," RFC 2396, T. Berners-Lee, R. Fielding,
532 L. Masinter, MIT/LCS, U.C. Irvine, Xerox Corporation, August 1998.
- 533 [4] WS-Message Delivery Version 1.0, <http://www.w3.org/Submission/2004/SUBM-ws->
534 [messagedelivery-20040426/](http://www.w3.org/Submission/2004/SUBM-ws-messagedelivery-20040426/)
- 535 [5] WS-Reliability latest specification, <http://www.oasis->
536 [open.org/committees/download.php/8909/WS-Reliability-2004-08-23.pdf](http://www.oasis-open.org/committees/download.php/8909/WS-Reliability-2004-08-23.pdf). See Section 4.2.3.2
537 (and its subsection), 4.3.1 (and its subsections). Please note that WS-R defines BareURI as the
538 default.
- 539 [6] Addressing wrapper schema, <http://www.oasis->
540 [open.org/apps/org/workgroup/wsrn/download.php/8365/reference-1.1.xsd](http://www.oasis-open.org/apps/org/workgroup/wsrn/download.php/8365/reference-1.1.xsd)
- 541 [7] WS-R schema that uses the serviceRefType, <http://www.oasis->
542 [open.org/apps/org/workgroup/wsrn/download.php/8477/ws-reliability-1.1.xsd](http://www.oasis-open.org/apps/org/workgroup/wsrn/download.php/8477/ws-reliability-1.1.xsd)
- 543 [8] Web Services Addressing, see <http://www.w3.org/Submission/ws-addressing/>
- 544 [9] Web Services Security: SOAP Message Security V1.0, <http://docs.oasis->
545 [open.org/wss/2004/01/oasis-200401-wss-soap-message-security-1.0.pdf](http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-soap-message-security-1.0.pdf)
- 546 [10] J. N. Gray, "The transaction concept: virtues and limitations", Proceedings of the 7th VLDB
547 Conference, September 1981, pp. 144-154.

548 **Appendix A. Acknowledgements**

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

551

Appendix B. Notices

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

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

564

565 **Copyright © OASIS Open 2005. All Rights Reserved.**

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

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

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

582

583

584