

1
1

2 REST-Atomic Transactions

3 32.0 draft 8

4

5 Version created 4 December 2012

6

7 **Editors**

8 Mark Little (mlittle@redhat.com)

9

10

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25

26

27

28

29

30

31

32

33

34

35

36

37

38 Abstract

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

48

49 Table of contents

501	Note on terminology.....	4
512	REST-Atomic Transaction.....	5
52	2.1 Relationship to HTTP.....	5
53	2.2 Header linking.....	5
54	2.3 The protocol.....	5
55	2.3.1 Two-phase commit.....	6
56	2.3.2 State transitions.....	7
57	2.3.3 Client and transaction interactions.....	8
58	2.3.3.1 Creating a transaction.....	8
59	2.3.3.2 Obtaining the transaction status.....	9
60	2.3.3.3 Terminating a transaction.....	10
61	2.3.4 Transaction context propagation.....	11
62	2.3.5 Coordinator and participant interactions.....	11
63	2.3.5.1 Enlisting a two-phase aware participant.....	11
64	2.3.5.2 Enlisting a two-phase unaware participant.....	12
65	2.3.5.3 Obtaining the status of a participant.....	13
66	2.3.5.4 Terminating a participant.....	13
67	2.3.6 Recovery.....	14
68	2.3.7 Pre- and post- two-phase commit processing.....	15
69	2.3.8 Statuses.....	16
703	Security Model.....	17
714	Security Considerations.....	18
725	References.....	19
73		

74 **1 Note on terminology**

75 The keywords "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD",
76 "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be
77 interpreted as described in RFC2119 [1].

78 Namespace URIs of the general form <http://example.org> and <http://example.com> represents
79 some application-dependent or context-dependent URI as defined in RFC 2396 [2].

80

81

82 **2 REST-Atomic Transaction**

83 *Atomic transactions* are a well-known technique for guaranteeing consistency in the presence of
84 failures [3]. The ACID properties of atomic transactions (Atomicity, Consistency, Isolation, and
85 Durability) ensure that even in complex business applications consistency of state is preserved,
86 despite concurrent accesses and failures. This is an extremely useful fault-tolerance technique,
87 especially when multiple, possibly remote, resources are involved.

88

89 Examples of coordinated outcomes include the classic two-phase commit protocol, a three phase
90 commit protocol, open nested transaction protocol, asynchronous messaging protocol, or
91 business process automation protocol. Coordinators can be participants of other coordinators.
92 When a coordinator registers itself with another coordinator, it can represent a series of local
93 activities and map a neutral transaction protocol onto a platform-specific transaction protocol.

94 **2.1 Relationship to HTTP**

95 This specification defines how to perform Atomic transactions using REST principles. However, in
96 order to provide a concrete mapping to a specific implementation, HTTP has been chosen.
97 Mappings to other protocols, such as JMS, is possible but outside the scope of this specification.

98 **2.2 Header linking**

99 Relationships between resources will be defined using the Link Header specification [4].

100 **2.3 The protocol**

101 The *REST-Atomic Transactions* model recognizes that HTTP is a good protocol for
102 interoperability as much as for the Internet. As such, interoperability of existing transaction
103 processing systems is an important consideration for this specification. Business-to-business
104 activities will typically involve back-end transaction processing systems either directly or indirectly
105 and being able to tie together these environments will be the key to the successful take-up of
106 Web Services transactions.

107

108 Although traditional atomic transactions may not be suitable for all Web based applications, they
109 are most definitely suitable for some, and particularly high-value interactions [such](#) as those
110 involved in finance. As a result, the Atomic Transaction model has been designed with
111 interoperability in mind. Within this model it is assumed that all services (and associated
112 participants) provide ACID semantics and that any use of atomic transactions occurs in
113 environments and situations where this is appropriate: in a trusted domain, over short durations.

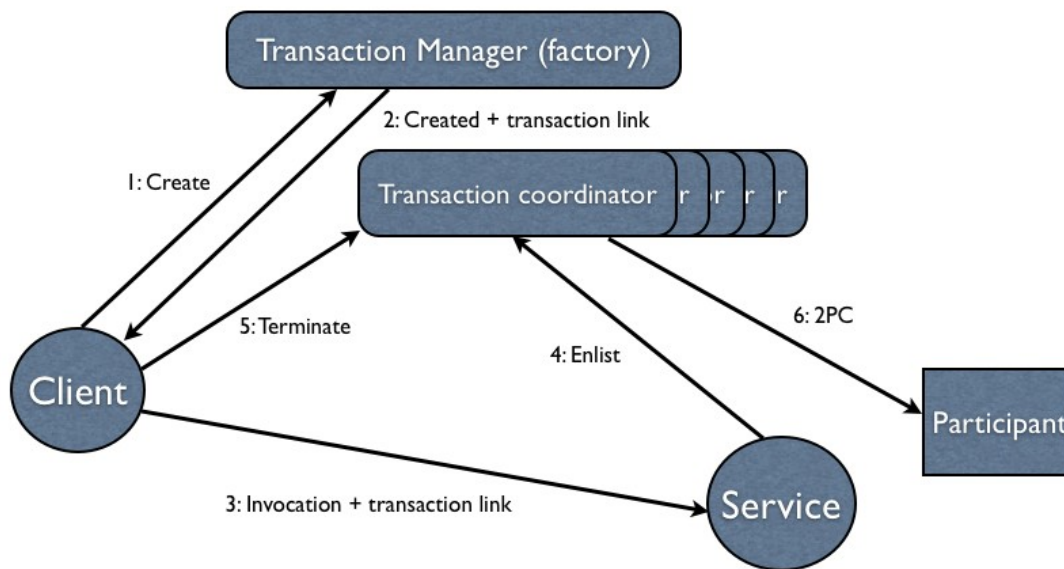
114

115 Note, this specification only defines how to accomplish atomic outcomes between participations
116 within the scope of the same transaction. It is assumed that if all ACID properties are required
117 then C, I and D are provided in some way outside this scope of this specification. This means that
118 some applications MAY use the REST-Atomic Transaction purely to achieve atomicity.

119

120 The following diagram illustrates the various components defined within this protocol. We shall
121 discuss each of these in the remainder of this specification.

122



123

1242.3.1 Two-phase commit

125The ACID transaction model uses a traditional two-phase commit protocol [3] with the following
126optimizations:

- 127 • *Presumed rollback*: the transaction coordinator need not record information about the
128 participants in stable storage until it decides to commit, i.e., until after the prepare phase
129 has completed successfully. A definitive answer that a transaction does not exist can be
130 used to infer that it rolled back.
- 131 • *One-phase*: if the coordinator discovers that only a single participant is registered then it
132 SHOULD omit the prepare phase.
- 133 • *Read-only*: a participant that is responsible for a service that did not modify any
134 transactional data during the course of the transaction can indicate to the coordinator
135 during prepare that it is a *read-only participant* and the coordinator SHOULD omit it from
136 the second phase of the commit protocol.

137

138Participants that have successfully passed the *prepare* phase are allowed to make autonomous
139decisions as to whether they commit or rollback. A participant that makes such an autonomous
140choice *must* record its decision in case it is eventually contacted to complete the original
141transaction. If the coordinator eventually informs the participant of the fate of the transaction and
142it is the same as the autonomous choice the participant made, then there is obviously no
143problem: the participant simply got there before the coordinator did. However, if the decision is
144contrary, then a non-atomic outcome has happened: a *heuristic outcome*, with a corresponding
145*heuristic decision*.

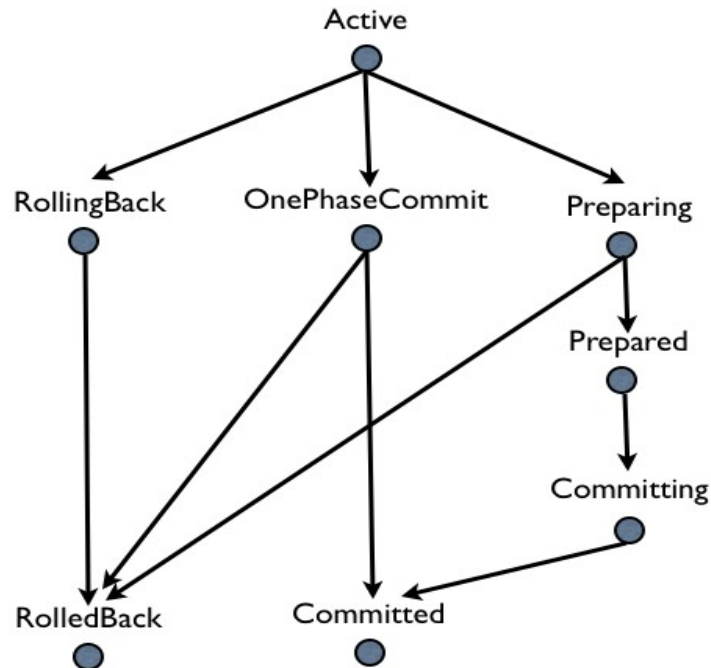
146

147The possible heuristic outcomes are:

- 148 • *Heuristic rollback*: the commit operation failed because some or all of the participants
149 unilaterally rolled back the transaction.
- 150 • *Heuristic commit*: an attempted rollback operation failed because all of the participants
151 unilaterally committed. This may happen if, for example, the coordinator was able to
152 successfully prepare the transaction but then decided to roll it back (e.g., it could not
153 update its log) but in the meanwhile the participants decided to commit.
- 154 • *Heuristic mixed*: some updates were committed while others were rolled back.
- 155 • *Heuristic hazard*: the disposition of some of the updates is unknown. For those which are
156 known, they have either all been committed or all rolled back.

1572.3.2 State transitions

158A transaction (coordinator and two-phase participant) goes through the state transitions shown
 159below. Note that non-atomic (heuristic) outcomes are not show on the diagram for simplicity, but
 160are discussed in a later section:



161

162There is a new media type to represent the status of a coordinator and its participants:
 163application/txstatus+json, which supports a return type based on the scheme maintained at
 164www.rest-star.org/... For example:

165tx-status+json=TransactionActive

166The EBNF definition of this media type is:

```

167
168<application/txstatus> ::= "tx-status+json" "=" <tx-state>
169 <tx-state> ::=
170     "TransactionRollbackOnly" |
171     "TransactionRollingBack" |
172     "TransactionRolledBack" |
173     "TransactionCommitting" |
174     "TransactionCommitted" |
175     "TransactionCommittedOnePhase"
176     "TransactionHeuristicRollback" |
177     "TransactionHeuristicCommit" |
178     "TransactionHeuristicHazard" |
179     "TransactionHeuristicMixed" |
180     "TransactionPreparing" |
181     "TransactionPrepared" |
182     "TransactionActive" |
183     "TransactionStatusUnknown"
  
```

184

185The text media type for a list of transactions (application/txlist) is simply a comma separated list
 186of transaction URLs. In EBNF:

187
 188 [transaction url list ::= url { "," url }*](#)
 189 [<url> ::= see RFC 1738](#)

190

191 2.3.3 Client and transaction interactions

192 The transaction [manager](#) is represented by a URI ([referred to as the transaction-manager URI](#)).
 193 In the rest of this specification we shall assume it is <http://www.fabrikam.com/transaction->
 194 [manager](#), but it could be any URI and its role need not be explicitly apparent within the structure
 195 of the URI.

196 2.3.3.1 Creating a transaction

197 Performing a POST on [/transaction-manager](#) [the transaction-manager URI](#) with [header](#) as shown
 198 below will start a new transaction with a default timeout. A successful invocation will return 201
 199 and the Location header MUST contain the URI of the newly created transaction resource, which
 200 we will refer to as transaction-coordinator in the rest of this specification. [At least two](#) related
 201 URLs MUST also be returned, one for the terminator of the transaction to use (typically referred
 202 to as the *client*) and one used for registering durable participation in the transaction (typically
 203 referred to as the *server*). [These are referred to as the transaction-terminator and transaction-](#)
 204 [enlistment URIs, respectively](#). Although uniform URL structures are used in the examples, these
 205 linked URLs can be of arbitrary format.

206
 207 [Note, an implementation MAY use the same URL for the terminator and participants.](#)

208
 209 POST /transaction-manager HTTP/1.1
 210 From: foo@bar.com

211
 212 The corresponding response would be:

213
 214 HTTP 1.1 201 Created
 215 Location: /transaction-coordinator/1234
 216 Link: </transaction-coordinator/1234/terminator>;
 217 rel="terminator",
 218 ~~Link:~~ </transaction-coordinator/1234/participant>;
 219 rel="durable--participant",
 220 ~~Link:~~ </transaction-coordinator/1234/vparticipant>;
 221 rel="volatile--participant"
 222

223 An implementation MAY return a Link reference for volatile participants if it supports the
 224 OPTIONAL volatile two-phase commit protocol, which is described later in this specification.

225
 226 [Note, the coordinator does not have to be co-located with the transaction manager resource, nor](#)
 227 [does it need to have the same URL prefix.](#)

228
 229 Performing a HEAD on [the-transaction-coordinator URI](#) MUST return the same link information.

230
 231 HEAD /transaction-coordinator/1234 HTTP/1.1
 232 From: foo@bar.com

233

234HTTP/1.1 200 OK
 235Link:</transaction-coordinator/1234/terminator>;
 236rel="terminator" ,
 237Link:</transaction-coordinator/1234/participant>;
 238rel="durable--participant" ,
 239Link:</transaction-coordinator/1234/vparticipant>;
 240rel="volatile--participant"
 241
 242Performing a POST on [the transaction-manager URI](#) as shown below will start a new transaction
 243with the specified timeout in milliseconds.
 244
 245POST /transaction-manager HTTP/1.1
 246From: foo@bar.com
 247Content-Type: text/plain
 248Content-Length: --
 249
 250timeout=1000
 251
 252If the transaction is terminated because of a timeout, the resources representing the created
 253transaction are deleted. All further invocations on the transaction-coordinator or any of its related
 254URIs MAY return 410 if the implementation records information about transactions that have
 255rolled back, (not necessary for presumed rollback semantics) but at a minimum MUST return 404.
 256The invoker can assume this was a rollback.
 257
 258A failure during the POST request, such as a network partition, may mean that the initial
 259response is not received. In this situation a client can retry the POST. Multiple transaction
 260coordinators may be created as a result, but the client SHOULD only use one of them and the
 261others will eventually timeout.
 262
 263Performing a GET on [the /transaction-manager URI with media type application/txlist](#) returns a
 264list of all transaction-coordinator URIs known to the coordinator (active and in recovery). [The](#)
 265[returned response MAY include a link header with rel attribute "statistics" linking to a resource](#)
 266[that contains statistical information such as the number of transactions that have committed and](#)
 267[aborted. The link MAY contain a media type hint with value "application/txstaturext+xml".](#)
 268
 269[Performing a GET on the transaction-manager URI with media type application/txstaturext+xml](#)
 270[returns extended information about the transaction-manager resource such as how long it has](#)
 271[been up and all transaction-coordinator URIs.](#)
 272

2732.3.3.2 Obtaining the transaction status

274Performing a GET on [the transaction-coordinator URI/transaction-coordinator/1234](#) returns the
 275current status of the transaction, as described later.
 276
 277GET /transaction-coordinator/1234 HTTP/1.1
 278Accept: application/txstatus
 279
 280With an example response:
 281
 282HTTP/1.1 200 OK

```

283Content-Length: --
284Content-Type: application/txstatus
285Link:</transaction-coordinator/1234/terminator>;
286rel="terminator",
287</transaction-coordinator/1234/participant>;
288rel="durable-participant",
289</transaction-coordinator/1234/vparticipant>;
290rel="volatile-participant"
291
292tx-status:txstatus=TransactionActive
293
294Performing a DELETE on any of the transaction-coordinator or transaction-enlistment URIs
295/transaction-coordinator URIs will return a 403.
296
297Additional information about the transaction, such as the number of participants and their
298individual URIs, MAY be returned if the client specifies the application/txstatus+xml and the
299implementation supports that type, otherwise status 415 is returned (as per RFC 2616)..

```

3002.3.3.3 Terminating a transaction

301The client can PUT one of the following to [the transaction-terminator URI](#) `/transaction-`
302`coordinator/1234/terminator` in order to control the outcome of the transaction; anything else
303MUST return a 400 (unless the terminator and transaction URLs are the same in which case GET
304would return the transaction status as described previously). Performing a PUT as shown below
305will trigger the commit of the transaction. Upon termination, the resource and all associated
306resources are implicitly deleted. For any subsequent PUT invocation, such as due to a
307timeout/retry, then an implementation MAY return 410 if the implementation records information
308about transactions that have rolled back, (not necessary for presumed rollback semantics) but at
309a minimum MUST return 404. The invoker can assume this was a rollback. In order for an
310interested party to know for sure the outcome of a transaction then it MUST be registered as a
311participant with the transaction coordinator.

```

312
313PUT /transaction-coordinator/1234/terminator HTTP/1.1
314From: foo@bar.com
315Content-Type: application/txstatus
316Content-Length: --
317
318tx-status:txstatus=TransactionCommitted

```

319
320The response body MAY contain the transaction outcome. If the transaction no longer exists then
321an implementation MAY return 410 if the implementation records information about transactions
322that have rolled back, (not necessary for presumed rollback semantics) but at a minimum MUST
323return 404.

324
325The state of the transaction MUST be TransactionActive for this operation to succeed. If the
326transaction is in an invalid state for the operation then the implementation MUST return a 412
327status code. Otherwise the implementation MAY return 200 or 202 codes. In the latter case the
328Location header SHOULD contain a URI upon which a GET may be performed to obtain the
329transaction outcome. It is implementation dependent as to how long this URI will remain valid.
330Once removed by an implementation then 410 MUST be returned.

331

332The transaction may be told to rollback with the following PUT request:

333

```
334PUT /transaction-coordinator/1234/terminator HTTP/1.1
```

```
335From: foo@bar.com
```

```
336Content-Type: application/txstatus
```

```
337Content-Length: --
```

338

```
339tx-statustxstatus=TransactionRolledBack
```

3402.3.4 Transaction context propagation

341When making an invocation on a resource that needs to participate in a transaction, either the
342transaction-coordinator URI or the [enlistingtransaction-enlistment](#) URI (e.g., /transaction-
343coordinator/1234/participant) needs to be transmitted to the resource. This specification does not
344mandate a mechanism for [propagation of this context information to the resource](#). However, the
345following OPTIONAL approaches are recommended.

346

347 ● The URI is passed as a Link with the relevant service interaction.

348 ● [Services participating in the transaction](#) return a Link to the client that can be used to
349 register participation with the coordinator.

350

351Note, a server SHOULD only use the URIs it is given directly and not attempt to infer any others.

3522.3.5 Coordinator and participant interactions

353Once a resource has the [transaction_or_enlistment](#) URI, it can register participation in the
354transaction. [Each participant must be uniquely identified to the transaction in order that the
355protocol can guarantee consistency and atomicity in the event of failure and recovery](#). The
356participant is free to use whatever URI structure it desires for uniquely identifying itself; in the rest
357of this specification we shall assume it is /participant-resource [and refer to it as the participant-
358resource URI](#).

3592.3.5.1 Enlisting a two-phase aware participant

360A participant is registered with [the #transaction-coordinator_](#) using POST on the [participant-Link-
361enlistment](#) URI obtained when the transaction was created originally. [The request must include
362two link headers: one to uniquely identify the participant to the coordinator and one to provide a
363terminator resource \(referred to as the participant-terminator URI\) that the coordinator will use to
364terminate the participant. If the rel attributes of the link are not participant and terminator
365the implementation must return 400. Note, the following URIs are only examples, and an
366implementation is free to use whatever structure/format it likes:](#)

367

```
368POST /transaction-coordinator/1234/participant HTTP/1.1
```

```
369From: foo@bar.com
```

```
370Link:</participant-resource>; rel="participant",
```

```
371</participant-resource/terminator>; rel="terminator"
```

372

```
373Content-Length: 0
```

374

375Performing a HEAD on [a registered participantthe participant-resource](#) URI MUST return the
376terminator reference, as shown below:

377
 378HEAD /participant-resource HTTP/1.1
 379From: foo@bar.com
 380
 381HTTP/1.1 200 OK
 382Link:</participant-resource/terminator>;
 383rel="terminator"

384
 385If the transaction is not TransactionActive [when registration is attempted](#), then the implementation
 386MUST return a 412 status code. If the implementation has seen this participant URI before then it
 387MUST return 400. Otherwise the operation is considered a success and the implementation
 388MUST return 201 and SHOULD use the Location header to give a participant specific URI that
 389the participant MAY use later during prepare or for recovery purposes. The lifetime of this URI is
 390the same as [the transaction-coordinator URI /transaction-coordinator](#). In the rest of this
 391specification we shall refer to this [URI](#) as [the participant-recovery URI /participant-recovery](#) (not
 392to be confused with the [/participant-resource URI](#)) ~~although the actual format is implementation-~~
 393~~dependent~~.

394
 395HTTP/1.1 201 Created
 396Location: /participant-recovery/1234
 397

3982.3.5.2 Enlisting a two-phase unaware participant

399In order for a participant to be enlisted with a transaction it MUST be transaction aware to fulfill
 400the requirements placed on it to ensure data consistency in the presence of failures or concurrent
 401access. However, it is not necessary that a participant be modified such that it has a terminator
 402resource as outlined previously: it simply needs a way to tell the coordinator which resource(s) to
 403communicate [with](#) when driving the two-phase protocol. This type of participant will be referred to
 404as Two-Phase Unaware, though strictly speaking such a participant or service does need to
 405understand the protocol as mentioned earlier.

406
 407Note, enlisting two-phase unaware participants is an OPTIONAL part of the specification. An
 408implementation that does not support this MUST return 405.

409
 410During enlistment a service MUST provide URIs for prepare, commit, rollback and OPTIONAL
 411commit-one-phase:

412
 413POST /transaction-coordinator/1234/participant HTTP/1.1
 414From: foo@bar.com
 415Link:</participant-resource>; rel="participant",
 416</participant-resource/prepare>; rel="prepare",
 417</participant-resource/commit>; rel="commit",
 418</participant-resource/rollback>; rel="rollback",
 419</participant-resource/commit-one-phase>; rel="commit-
 420one-phase"

421
 422Content-Length: 0

423
 424Performing a HEAD on a registered participant URI MUST return these references, as shown
 425below:

426
 427HEAD /participant-resource HTTP/1.1
 428From: foo@bar.com
 429
 430HTTP/1.1 200 OK
 431Link:</participant-resource/prepare>; rel="prepare",
 432Link:</participant-resource/commit>; rel="commit",
 433Link:</participant-resource/rollback>; rel="rollback",
 434</participant-resource/commit-one-phase>; rel="commit-
 435one-phase"
 436
 437
 438A service that registers a participant MUST therefore either define a *terminator* relationship for
 439the participant or the relationships/resources needed for the two-phase commit protocol.

4402.3.5.3 Obtaining the status of a participant

441Performing a GET on the */participant-resource* [URI](#) MUST return the current status of the
 442participant in the same way as for the */transaction-coordinator* URI discussed earlier. Determining
 443the status of a participant whose URI has been removed is similar to that discussed for [the](#)
 444*/transaction-coordinator* [URI](#).

445
 446GET /participant-resource/1234 HTTP/1.1
 447Accept: application/txstatus
 448
 449With an example response:
 450
 451HTTP/1.1 200 OK
 452Content-Length: --
 453Content-Type: application/txstatus
 454
 455~~tx-status~~txstatus=TransactionActive

4562.3.5.4 Terminating a participant

457The coordinator drives the participant through the two-phase commit protocol by sending a PUT
 458request to the participant terminator URI provided during enlistment, with [the desired transaction](#)
 459[outcome as the content](#) (TransactionPrepared, TransactionCommitted, TransactionRolledBack or
 460TransactionCommittedOnePhase). [For instance, here is how the prepare phase would be driven:](#)

461
 462PUT /participant-resource/terminator HTTP/1.1
 463From: foo@bar.com
 464Content-Type: application/txstatus
 465Content-Length: --
 466
 467~~tx-status~~txstatus=TransactionPrepared

468
 469If [PUT](#) is successful then the implementation MUST return 200. A subsequent GET on the URI
 470will return the current status of the participant as described previously. It is not always necessary

471to enquire as to the status of the participant once the operation has been successful.

472

473If [PUT](#) fails, e.g., the participant cannot be prepared, then the implementation MUST return 409.

474It is implementation [dependent](#) as to whether the `/participant-resource` or related URIs

475remain valid, i.e., an implementation MAY delete the resource as a result of a failure. Depending

476upon the point in the two-phase commit protocol where such a failure occurs the transaction

477MUST be rolled back, [e.g., because we use presumed abort semantics, failures prior to the end](#)

478of the [prepare phase MUST result in a roll back](#). If the participant is not in the correct state for the

479requested operation, e.g., [TransactionPrepared](#) when it has been already been prepared, then

480the implementation MUST return 412.

481

482If the transaction coordinator receives any response other than 200 for Prepare then the

483transaction MUST rollback.

484

485After [a request to change the resource state using](#) [TransactionRolledBack](#),

486[TransactionCommitted](#) or [TransactionCommittedOnePhase](#), [any](#) subsequent [PUT](#) request MUST

487return a [409](#) or 410 code.

488

489Note, read-only MAY be modeled as a DELETE request from the participant to the coordinator

490using the URI returned during registration in the Location header, as mentioned previously, i.e.,

491[the /participant-recovery URI](#). If GET is used to obtain the status of the participant after a 200

492response is received to the original PUT for Prepare then the implementation MUST return 410 if

493the participant was read-only.

494

495The usual rules of heuristic decisions apply here (i.e., the participant cannot forget the choice it

496made until it is told to by the coordinator).

497

498Performing a [DELETE](#) on [the /participant-resource URI](#) will cause the participant to forget any

499heuristic decision it made on behalf of the transaction. If the operation succeeds then 200 MUST

500be returned and the implementation MAY delete the resource; a subsequent PUT or GET request

501MUST return 410. Any other response means the coordinator MUST retry.

502.3.6 Recovery

503In general it is assumed that failed actors in this protocol, i.e., coordinator or participants, will

504recover on the same URI as they had prior to the failure. [HTTP provides a number of options to](#)

505[support temporary or permanent changes of address, including 301 \(Moved Permanently\) and](#)

506[307 \(Temporary Redirect\)](#). If that is not possible then these endpoints SHOULD return a 301

507status code or some other way of indicating that the participant has moved elsewhere. HTTP

508response codes such as 307 MAY also be used by the implementation if a temporary redirection

509is used.

510

511However, sometimes it is possible that a participant may crash and recover on a different URI,

512e.g., the original machine is unavailable, or that for expediency it is necessary to move recovery

513to a different machine. In that case it may be that transaction coordinator is unable to complete

514the transaction, even during recovery. As a result this protocol defines a way for a recovering

515server to update the information maintained by the coordinator on behalf of these participants.

516

517If the recovering participant uses the `/participant-recovery` URI returned by the coordinator during

518enlistment then a GET on [the /participant-recovery URI](#) will return [the participant resource and](#)

519[terminator as link headers the original participant URI supplied when that the the participant was](#)

520[registered used during the original registration](#).

521

522Performing a PUT on [the /participant-recovery URI](#) will overwrite the old participant URI with the

523new one supplied. [This operation is equivalent to re-enlisting the participant](#). This will also trigger

524off a recovery attempt on the associated transaction using the new participant URI. [For example](#)

525to update location URIs, a two phase aware participant would PUT the following document:

526

527PUT /participant-recovery/1234 HTTP/1.1

528From: **foo@bar.com**

529Link:</new-participant-resource>; rel="participant",

530</participant-resource/new-terminator>;

531rel="terminator"

532Content-Type: text/plain

533Content-Length: --0

534

535new-address=URI

536

537Similarly for a two phase unaware participant.

538

539If, after performing the PUT request to the participant-recovery URI, the participant is not asked to
540complete (within an implementation dependent period) then it SHOULD reissue the PUT request.

5412.3.7 Pre- and post- two-phase commit processing

542Most modern transaction processing systems allow the creation of participants that do not take
543part in the two-phase commit protocol, but are informed before it begins and after it has
544completed. They are called *Synchronizations*, and are typically employed to flush volatile
545(cached) state, which may be being used to improve performance of an application, to a
546recoverable object or database prior to the transaction committing.

547

548This additional protocol is accomplished in this specification by supporting an additional two-
549phase commit protocol that enclosed the protocol we have already discussed. This will be termed
550the Volatile Two Phase Commit protocol, as the participants involved in it are not required to be
551durable for the purposes of data consistency, whereas the other protocol will be termed the
552Durable Two Phase Commit protocol. The coordinator MUST not record any durable information
553on behalf of Volatile participants.

554

555In this case the Volatile prepare phase executes prior to the Durable prepare where the
556transaction-coordinator sends a PUT request to the registered volatile-participant: only if this
557prepare succeeds will the Durable protocol be executed. The volatile-participant MUST indicate
558success by returning a 200 status code (any other code indicates failure). If the Durable protocol
559completes then this MAY be communicated to the Volatile participants through the commit or
560rollback phases. In this case the transaction-coordinator sends a PUT request to the registered
561volatile-participant with with the outcome in the request body (using content type
562application/txstatus). However, because the coordinator does not maintain any information about
563these participants and the Durable protocol has completed, this SHOULD be a best-effort
564approach only, i.e., such participants SHOULD NOT assume they will be informed about the
565transaction outcome. If that is a necessity then they should register with the Durable protocol
566instead.

567

568The Volatile protocol is identical to the Durable protocol described already. The only differences
569are as discussed below:

570

- 571 • It is an OPTIONAL protocol. An implementation that supports the protocol MUST show this
- 572 when the transaction is created through a Link relationship: it returns an additional Linked
- 573 resource whose relationship is defined as "volatile-participant". Services MUST use this
- 574 URI when registering volatile participants.

- 575 • There is no recovery associated with the Volatile protocol. Therefore the /participant-
576 recovery URI SHOULD NOT be used by an implementation.
- 577 • There can be no heuristic outcomes associated with the Volatile protocol.
- 578 • An implementation MAY allow registration in the Volatile protocol after the transaction has
579 been asked to terminate as long as the Durable protocol has not started.
- 580 • There is no one-phase commit optimization for the Volatile protocol.

5812.3.8 Statuses

582 [Resources MUST](#) return the following statuses by performing a GET on the appropriate
583/transaction-coordinator or participant URI:

- 584 • TransactionRollbackOnly: the status of the endpoint is that it will roll back eventually.
- 585 • TransactionRollingBack: the endpoint is in the process of rolling back. [If the recipient has
586 already rolled back then it MUST return a 410 error code.](#)
- 587 • TransactionRolledBack: the endpoint has rolled back.
- 588 • TransactionCommitting: the endpoint is in the process of committing. This does not mean
589 that the final outcome will be Committed. [If the recipient has already committed then it
590 MUST return a 410 error code.](#)
- 591 • TransactionCommitted: the endpoint has committed.
- 592 • [TransactionCommittedOnePhase: the recipient has committed the transaction without
593 going through a prepare phase. If the recipient has previously been asked to prepare
594 then it MUST return a 412 error code. If the recipient has already terminated, then it
595 MUST return a 410 error code.](#)
- 596 • TransactionHeuristicRollback: all of the participants rolled back when they were asked to
597 commit.
- 598 • TransactionHeuristicCommit: all of the participants committed when they were asked to
599 rollback.
- 600 • TransactionHeuristicHazard: some of the participants rolled back, some committed and the
601 outcome of others is indeterminate.
- 602 • TransactionHeuristicMixed: some of the participants rolled back whereas the remainder
603 committed.
- 604 • TransactionPreparing: the endpoint is preparing.
- 605 • TransactionPrepared: the endpoint has prepared.
- 606 • TransactionActive: the transaction is active, i.e., has not begun to terminate.
- 607 • [TransactionStatusUnknown: the status of the transaction is unknown](#)

608
609 [The statuses are also used to drive the two-phase commit protocol as discussed previously.](#)

6103 Security Model

611The security model for atomic transactions builds on the standard HTTP security model. That is,
612services have policies specifying their requirements and requestors provide claims (either implicit
613or explicit) and the requisite proof of those claims. Coordination context creation establishes a
614base secret which can be delegated by the creator as appropriate.

615

616Because atomic transactions represent a specific use case rather than the general nature of
617coordination contexts, additional aspects of the security model can be specified.

618

619All access to atomic transaction protocol instances is on the basis of identity. The nature of
620transactions, specifically the uncertainty of systems means that the security context established
621to register for the protocol instance may not be available for the entire duration of the protocol.
622Consider for example the scenarios where a participant has committed its part of the transaction,
623but for some reason the coordinator never receives acknowledgement of the commit. The result
624is that when communication is re-established in the future, the coordinator will attempt to confirm
625the commit status of the participant, but the participant, having committed the transaction and
626forgotten all information associated with it, no longer has access to the special keys associated
627with the token.

628

629There are, of course, techniques to mitigate this situation but such options will not always be
630successful. Consequently, when dealing with atomic transactions, it is critical that identity claims
631always be proven to ensure that coordinators maintain correct access control.

632

633There is still value in coordination context-specific tokens because they offer a bootstrap
634mechanism so that all participants need not be pre-authorized. As well, it provides additional
635security because only those instances of an identity with access to the token will be able to
636securely interact with the coordinator (limiting privileges strategy).

637

638The "list" of authorized participants ensures that application messages having a coordination
639context are properly authorized since altering the coordination context ID will not provide
640additional access unless (1) the bootstrap key is provided, or (2) the requestor is on the
641authorized participant "list" of identities.

642 **4 Security Considerations**

643 It is strongly RECOMMENDED that the communication between services be secured using HTTP
 644 security mechanisms. In order to properly secure messages, the body and all relevant headers
 645 need to be included in the signature. In the event that a participant communicates frequently with
 646 a coordinator, it is RECOMMENDED that a security context be established
 647.

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

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

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

660

661 Note, the content of Link header fields is not secure, private or integrity-guaranteed, and due
 662 caution should be exercised when using it. Use of Transport Layer Security (TLS) with HTTP [5]
 663 and [6] is currently the only end-to-end way to provide such protection.

664 **5** References

- 665[1] "Key words for use in RFCs to Indicate Requirement Levels," RFC 2119, S. Bradner, Harvard
666University, March 1997.
- 667[2] "Uniform Resource Identifiers (URI): Generic Syntax," RFC 2396, T. Berners-Lee, R. Fielding,
668L. Masinter, MIT/LCS, U.C. Irvine, Xerox Corporation, August 1998.
- 669[3] J. N. Gray, "The transaction concept: virtues and limitations", Proceedings of the 7th VLDB
670Conference, September 1981, pp. 144-154.
- 671[4] M. Nottingham, "HTTP Header Linking", [http://www.mnot.net/drafts/draft-nottingham-http-link-
672header-07.txt](http://www.mnot.net/drafts/draft-nottingham-http-link-672header-07.txt), June 2006.
- 673[5] <http://tools.ietf.org/html/rfc2818>
- 674[6] <http://tools.ietf.org/html/rfc2817>