Web Services ACID Specification (WS-ACID)

Editors draft version 0.2

Deleted: 1

Version created 8 July 2005

Deleted: 30 June 2005

Editors

Mark Little (mark.little@arjuna.com)

Eric Newcomer (eric.newcomer@iona.com)

Greg Pavlik (greg.pavlik@oracle.com)

Abstract

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for recovery and compensation, because machines may fail, processes may be cancelled, or services may be moved or withdrawn.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

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

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

Table of contents

Deleted: 1 Note on terminology 3¶
1.1 1.1 Namespace 3¶
1.1.1 1.1.1 Prefix Namespace 3¶
1.2 Referencing Specifications 3¶
2 Introduction 4¶
2.1 Problem statement 4¶
3 Architecture 6¶
3.1 Invocation of Service Operations 6¶
3.2 Relationship to WSDL 7¶
3.3 Referencing and addressing conventions 7¶
4 WS-ACID 9¶
4.1 Interposition 9¶
4.2 Restrictions imposed on using WS-CF 9¶
4.3 Two-phase commit 9¶
4.3.1 Coordinator state transitions for two-phase commit protocol 10¶
4.3.2 Two-phase participant state transitions 12¶
4.3.3 Two-phase commit message interactions 13¶
4.3.4 Pre- and post- two-phase commit processing 15¶
4.3.5 Coordinator state transitions for synchronization protocol 17¶
4.3.6 Recovery and interposition 17¶
4.3.7 The context 18¶
4.3.8 Statuses 18¶
5 References 20¶

1 Note on terminology

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

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

1.1 Namespace

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

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

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/wsrm/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**

1.2 Referencing Specifications

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

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

1.3 Precedence of schema and WSDL

Throughout this specification, WSDL and schema elements may be used for illustrative or convenience purposes. However, in a situation where those elements within this document differ from the separate WS-Context WSDL or schema files, it is those files that have precedence and not this specification.

Deleted: The keywords "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC2119 [2].¶ Namespace URIs of the general form "some-URI" represents some application-dependent or context-dependent URI as defined in RFC 2396 [3].¶

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

Formatted: Heading 2,H2

2 Architecture

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.

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.

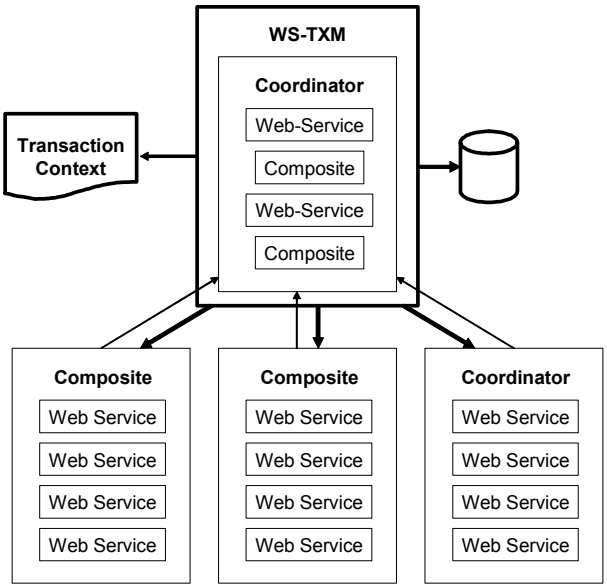

Figure 1, Relationship of transactions to coordination framework.

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.

2.1 Invocation of Service Operations

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

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.

Deleted: Introduction

Deleted: ¶
The concepts of atomic transactions have played a cornerstone role in creating today's enterprise application environments by providing guaranteed consistent outcome in complex multiparty business operations and a useful separation of concerns in applications. While numerous multiparty business applications involve various patterns based on atomic transactions in order to solve non-trivial business problems, it was not until recently the word "business transactions" accumulated any concrete meaning. Rapid developments in Internet infrastructure and protocols have yielded a new type of application interoperation concept that makes concepts which could only previously be considered in an abstract form an implementation reality. The effects of such changes have been felt most strongly in business environments, fuelling the mindset for a transition from traditional atomic transactions to extended transaction models better suited for Internet interoperation.¶
Most business-to-business applications require transactional support in order to guarantee consistent outcome and correct execution. These applications often involve long running computations, loosely coupled systems and components that do not share data, location, or administration and it is thus difficult to incorporate traditional ACID transactions within such architectures. For example, an airline reservation system may reserve a seat on a flight for an individual for a specific period of time, but if the individual does not confirm the seat within that period it will be unreserved.¶
The structuring mechanisms available within traditional transaction systems are sequential and concurrent composition of transactions. These mechanisms are sufficient if an applicatio ... [1]

Formatted: Bullets and Numbering

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

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

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

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

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

2.2 Relationship to WSDL

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

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

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

2.3 Referencing and addressing conventions

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

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

Deleted: How application services are invoked is outside the scope of this specification; however, context information related to the sender's activity needs to be referenced and/or propagated. ¶
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.¶
In order to support both synchronous and asynchronous interactions, the components are described in terms of the behavior and the interactions that occur between them. All interactions are described in terms of correlated messages, which a referencing specification MAY abstract at a higher level into request/response pairs. ¶
Faults and errors that may occur when a service is invoked are communicated back to other Web services in the activity via SOAP messages that are part of the standard protocol. The fault mechanism of the underlying SOAP-based transport isn't used. For example, if an operation fails because no activity is present when one is required, then it will be valid for the InvalidContextFault message to be received by the response service. To accommodate other errors or faults, all response service signatures have a generalFault operation as well as a transientFault operation.¶
Note, a transientFault message is produced when the implementation finds it cannot successfully execute the requested operation at that time from some *temporary* reason. This reason may be implementation or referencing specification specific. A receiver of a transientFault is free to retry the operation which originally generated it on the assumption that eve ... [2]

Formatted: Bullets and Numbering

Deleted: Coordination Framework

Formatted: Bullets and Numbering

```
<xsd:complexType name="ServiceRefType">
    <xsd:sequence>
      <xsd:any namespace="##other" processContents="lax"/>
    </xsd:sequence>
    <xsd:attribute name="reference-scheme" type="xsd:anyURI"
        use="optional"/>
  </xsd:complexType>
```

Figure 2, service-ref Element

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

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

Figure 3, ServiceRefType example.

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

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

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

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

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

Deleted: <xsd:schema targetNamespace="http://docs.oasis-open.org/wsrm/2004/06/reference-1.1.xsd" xmlns:xsd="http://www.w3.org/2001/XMLSchema" elementFormDefault="qualified" attributeFormDefault="unqualified" version="1.1">¶

<xsd:complexType name="ServiceRefType">¶

<xsd:sequence>¶

<xsd:any namespace="##other" processContents="lax" /> ¶

</xsd:sequence>¶

<xsd:attribute name="reference-scheme" type="xsd:anyURI" use="optional" /> ¶

</xsd:complexType>¶

Deleted: 1

Deleted: 2

Formatted: Font: Bold

3 WS-ACID

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

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

In the ACID model, each activity is bound to the scope of a transaction, such that the end of an activity automatically triggers the termination (commit or rollback) of the associated transaction. The coordinator-type URI for the ACID transaction model is http://www.webservicestransactions.org/wsdl/wstxm/tx-acid/2003/03

3.1 Restrictions imposed on using WS-CF

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

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

3.2 Two-phase commit

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

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

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

The possible heuristic outcomes are:

Formatted: Bullets and Numbering

Deleted: defined in WS-TXM

Comment: Update.

Deleted: <#>Interposition¶ WS-CF supports the notion of *interposition*: where a Participant Service that is enlisted with a Registration Service also behaves as a Registration Service to other Participant Services. In this way, WS-CF supports the building of graphs and trees by the addition of participants to an activity structure that are themselves registration endpoints.¶

Formatted: Bullets and Numbering

Deleted: As well as the restrictions outlined previously for general WS-TXM protocols, the

Deleted: additional

Deleted: sent by the coordinator

Deleted: <#>It is illegal to call the WS-CF *coordinate* operation on the coordinator. All WS-TXM coordinator implementations will return the *notCoordinated* message if the coordinator receives a *coordinate* request.¶

Formatted: Bullets and Numbering

Deleted: there is only a single participant involved in the transaction then there is no need for a prepare phase since consensus is implicit

Deleted: it

Deleted: will be

Deleted: ted

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

3.2.1 WS-ACID and WS-Context message interactions

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

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

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

A transactional activity can either be committed or rolled back via the **wsctx:complete** message. The protocol-specific termination extension within the **wsctx:complete** message contains either the **wsacid:Commit** or **wsacid:Rollback** completion code, depending upon whether the transaction is to be committed or rolled back respectively.

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

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

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

Deleted: Coordinator s

Deleted: tate transitions for two-phase commit protocol

Formatted: Bullets and Numbering

Formatted: Font: Italic

Deleted: As shown in Figure 5, and in line with the basic WS-CTX, w

Deleted: the

Formatted: Font: Bold

Formatted: Note

Deleted: activity begins a *begin* message is sent by the activity service to the ACID transaction protocol ALS (Tx ALS); this then creates a corresponding coordinator that is associated with the activity through the context.

Deleted: begins in the *Active* state and

Deleted: . What this means is that if

Deleted: and as a result the Context Service terminates the activity, the

Deleted: also

Deleted: same state as the activity

Formatted: Font: Bold

Deleted: <#>¶

Deleted: activity

Deleted: complete in the Success state

Deleted: activity service

Formatted: Font: Bold

Deleted: *completeWithStatus* message to the Tx ALS which will then try to commit

Deleted: there is no need for the coordinator to execute the two-phase protocol

Deleted: **The activity completion status is either *Failure* or *Success*** ... [3]

Formatted: Font: Bold

Formatted: Font: Bold

Deleted: ; the activity ... [4]

Formatted: Font: Bold

Deleted: other

Deleted: , with the activ ... [5]

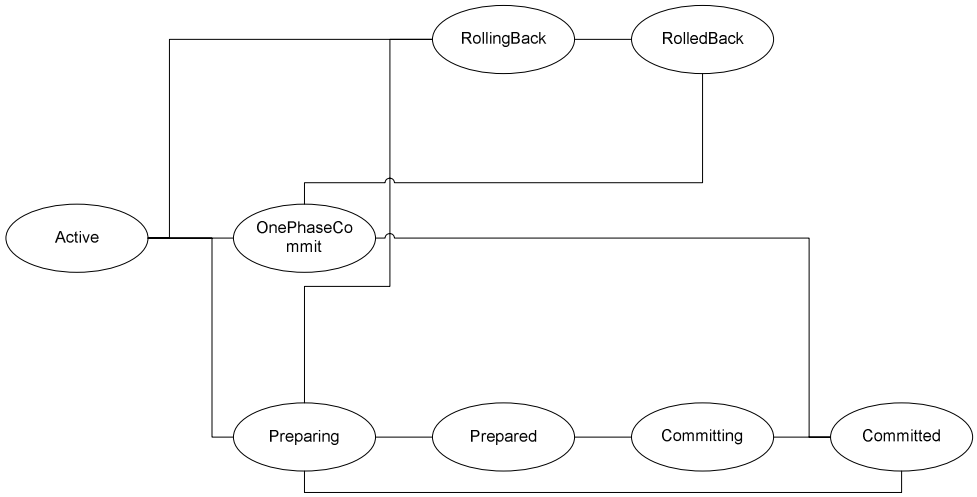

Figure 4, Transaction coordinator two-phase status transition.

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.

If the transaction commits then the protocol-specific termination status within the **wsctx:completed** message will contain the **wsacid:Committed** code. If the transaction rolls back then the code will be **wsacid:RolledBack**. All other errors are communicated using the fault model described earlier.

3.2.2 Coordinator and participant message interactions

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.

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.

The two-phase commit sub-protocol URI is http://www.webservicestransactions.org/wsdl/wstxm/tx-acid/2pc/2003/03 and this is used in the **wscf:addParticipant** message. If the OPTIONAL unique endpoint reference is returned in the **wscf:participantAdded** message then the participant MUST use this for sending coordination signals unless the addressing implementation dictates otherwise.

An enlisted Participant Service should expect to receive the following messages (illustrated in Figure 5):

- 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**

Deleted: 5

Formatted: Font: Bold

Deleted: In terms of the underlying activity service and coordination service, Figure 6 shows the flow of messages:¶ <#>The application starts a new activity, which ultimately causes the issuing of a *begin* on the Transaction/Coordination Service ALS to demarcate the beginning of a transaction. This causes the Transaction/Coordination Service to create a coordinator used to identify the activity instance and subsequently track the elements interested in the transaction outcome (i.e., the participants).¶ <#>The application issues a server method. Context information is appended to the message. The context is used at the target to recreate the execution environment. Having been passed a coordina ... [6]

Formatted: Bullets and Numbering

Deleted: that are

Deleted: because WS-CF supports an asynchrono ... [7]

Deleted: ,

Deleted: may

Deleted: “

Deleted: ”

Deleted: via the *setResponse* message

Deleted: implicitly

Deleted: AT

Deleted: for the two-phase commit protocol

Comment: Need to update.

Formatted: Font: Bold

Formatted: Font: Bold, Not Italic

Deleted: invocation

Formatted: Font: Bold

Deleted: The

Deleted: accepts

Deleted: Figure 8

Deleted: . The CoordinatorParticipant e ... [8]

Formatted: Font: Bold

protocol. The **wsacid:voteReadonly** and wsacid: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:HeuristicHazard** error code will be returned. Alternatively, if a subordinate coordinator finds that some of the participants have committed and some have rolled back then it must return the **wsacid:HeuristicMixed** error message.

- 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:HeuristicHazard** fault code. If the participant is a subordinate coordinator and some of its enlisted participants committed then it must return the **wsacid:HeuristicMixed** fault code. If the participant commits rather than rolls back then it must return the **wsacid:HeuristicCommit** 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:HeuristicHazard** fault code. If the participant is a subordinate coordinator and some of its enlisted participants committed then it must return the **wsacid:HeuristicMixed** fault code. If the participant rolls back rather than commits then it must return the **wsacid:HeuristicRollback** 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:HeuristicHazard** fault code. If the participant is a subordinate coordinator and some of its enlisted participants committed then it must return the **wsacid:HeuristicMixed** 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.

Comment: Issue 275
Formatted ... [9]
Deleted: it…must retu ... [10]
Formatted: Font: Bold
Deleted: h
Formatted: Font: Bold, Not Italic
Deleted: Fault…mess ... [11]
Comment: SOAP faults
Formatted: Font: Bold
Deleted: h
Formatted: Font: Bold, Not Italic
Deleted: Fault
Formatted: Font: Bold
Deleted: message
Deleted: cancelling
Formatted: Font: Bold
Deleted: the
Formatted: Font: Bold
Deleted: h
Formatted: Font: Bold, Not Italic
Deleted: Fault…mess ... [12]
Formatted: Font: Bold
Deleted: h
Formatted: Font: Bold, Not Italic
Deleted: Fault…mess ... [13]
Formatted: Font: Bold
Deleted: h
Formatted: Font: Bold, Not Italic
Deleted: Fault
Formatted ... [14]
Deleted: confirming…If the participant is a subordinate coordinator and cannot determine how all of its enlisted participants ter ... [15]
Formatted: Font: Bold
Deleted: h
Formatted ... [16]
Deleted: Fault…mess ... [17]
Deleted: it is possible ... [18]
Formatted ... [19]

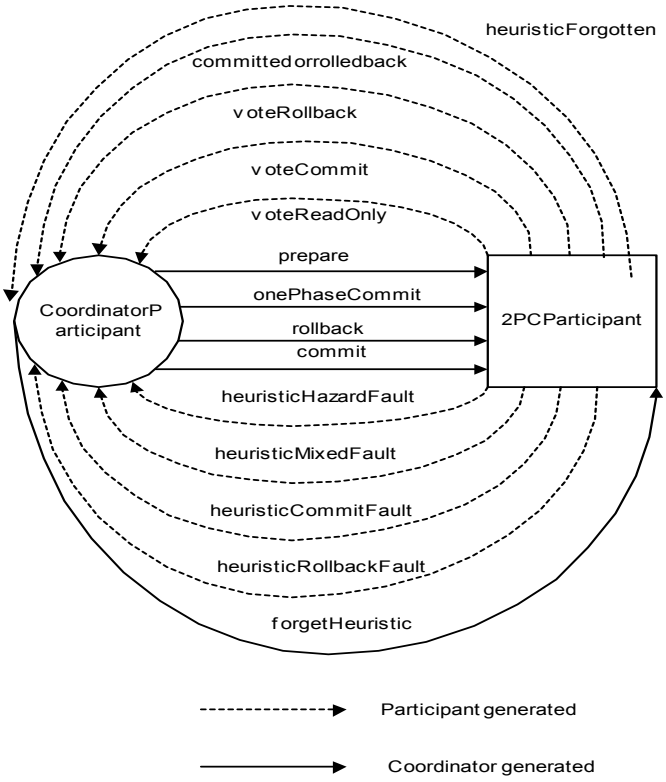

Figure 5, Coordinator-to-participant message exchanges.

The WSDL portType declarations for the CoordinatorParticipant and twoPCParticipant roles are shown in Figure 6.

```
<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">
```

Deleted: 8

Deleted: AT c

Comment: Need to change the names.

Deleted: Figure 9

Formatted: Code

```
    <wsdl:input message="tns:VoteMessage"/>
  </wsdl:operation>
  <wsdl:operation name="heuristicForgotten">
    <wsdl:input message="tns:HeuristicForgottenMessage"/>
  </wsdl:operation>
  <wsdl:operation name="heuristicFault">
    <wsdl:input message="tns:HeuristicFaultMessage"/>
  </wsdl:operation>
</wsdl:portType>
```

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

3.3 Pre- and post- two-phase commit processing

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

The sub-protocol URI for the synchronization protocol is http://www.webservicestransactions.org/wsdl/wstxm/tx-acid/sync/2003/03 and this is used in the **wscf:addParticipant** invocation.

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

- beforeCompletion: A Synchronization participant is informed that the coordinator it is registered with is about to complete the two-phase protocol and in what state, i.e., committing or rolling back. Any failure at this stage will cause the coordinator to rollback if it is not already doing so. Success is indicated by the **wsacid:beforeCompleted** message.
- afterCompletion: A Synchronization participant is informed that the coordinator it is registered with has completed the two-phase protocol and in what state, i.e., committed or rolled back (via the associated **wsacid:status**). Any failures at this stage have no affect on the transaction: an implementation MAY report these failures. Success is indicated by the **wsacid:afterCompleted** message.

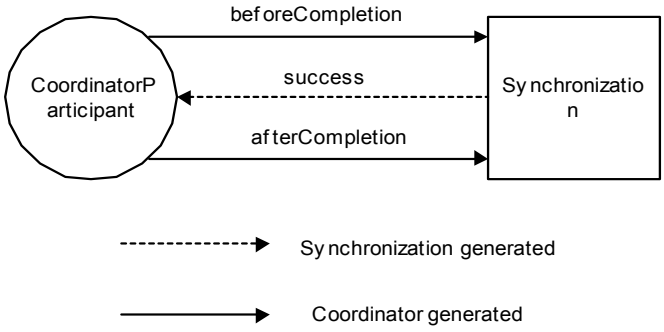

Figure 7, AT coordinator-to-synchronization message exchanges.

The WSDL portType declarations for the CoordinatorParticipant and Synchronization roles are shown in Figure 8.

```
<wsdl:portType name="SynchronizationPortType">
  <wsdl:operation name="beforeCompletion">
```

Deleted: 9

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

Formatted: Heading 2,H2

Formatted: Bullets and Numbering

Formatted: Font: Italic

Deleted: AT

Comment: Needs updating.

Formatted: Font: Bold

Formatted: Font: Bold, Not Italic

Deleted: Figure 10

Deleted: The failure of the participant

Deleted: cancel

Formatted: Font: Bold

Comment: Ensure consistency.

Formatted: Font: Bold

Deleted: Status

Formatted: Font: Bold

Deleted: by the participant

Formatted: Font: Bold

Deleted: 10

Deleted: Figure 11

Formatted: Code

```
    <wsdl:input message="tns:BeforeCompletionMessage"/>
  </wsdl:operation>
  <wsdl:operation name="afterCompletion">
    <wsdl:input message="tns:AfterCompletionMessage"/>
  </wsdl:operation>
</wsdl:portType>
<wsdl:portType name="CoordinatorParticipantPortType">
  <wsdl:operation name="beforeCompletionParticipantRegistered">
    <wsdl:input
message="tns:BeforeCompletionParticipantRegisteredMessage"/>
  </wsdl:operation>
  <wsdl:operation name="afterCompletionParticipantRegistered">
    <wsdl:input
message="tns:AfterCompletionParticipantRegisteredMessage"/>
  </wsdl:operation>
</wsdl:portType>
```

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

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

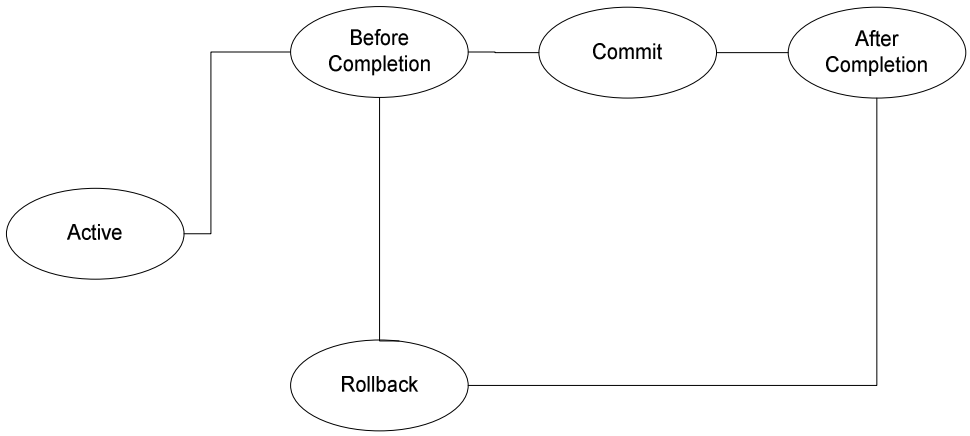

Figure 9, Transaction coordinator Synchronization state transitions.

The coordinator moves into the *BeforeCompletion* state and sends each enrolled Synchronization the **wsacid:beforeCompletion** message. Any error received by the coordinator from a Synchronization at this stage will force the transaction to roll back. Assuming no errors occur, the two-phase commit protocol is executed, as detailed previously. Once the protocol has completed, the coordinator transits to the *AfterCompletion* status and sends the **wsacid:afterCompletion** message to all Synchronizations; any errors at this stage do not affect the transaction outcome and how they are dealt with is implementation dependant.

Deleted: *11*

Deleted: Note, the participant is registered for both beforeCompletion and afterCompletion.¶
<#>Coordinator state transitions for synchronization protocol¶

Deleted: transitions

Deleted: 12

Formatted: Font: Bold

Formatted: Font: Bold, Not Italic

Formatted: Font: Bold

Formatted: Font: Bold, Not Italic

3.4 Checked transactions

Checked transactions have a number of integrity constraints including:

- Ensuring that only the transaction originator can commit the transaction.
- Ensuring that a transaction will not commit until all transactional invocations involved in the transaction have completed.

Some implementations will enforce checked behavior for the transactions they support, to provide an extra level of transaction integrity. The purpose of the checks is to ensure that all transactional requests made by the application have completed their processing before the transaction is committed. A checked Transaction Service guarantees that commit will not succeed unless all invocations involved in the transaction have completed. Rolling back the transaction does not require such as check, since all outstanding transactional activities will eventually rollback if they are not told to commit

There are many possible implementations of checked transactions. One provides equivalent function to that provided by the request/response inter-process communication models defined by X/Open. It describes the transaction integrity guarantees provided by many existing transaction systems. In X/Open, completion of the processing of a request means that the service has completed execution of its invocation and replied to the request. The level of transaction integrity provided by a Transaction Service implementing the X/Open model of checking provides equivalent function to that provided by the XATMI and TxRPC interfaces defined by X/Open for transactional applications.

3.5 Recovery and interposition

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

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

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

3.6 The context

```
<xs:complexType name="ContextType">
  <xs:complexContent>
    <xs:extension base="wstxm:ContextType"/>
  </xs:complexContent>
</xs:complexType>
<xs:element name="context" type="tns:ContextType"/>
```

Figure 10, Transaction Context.

3.7 Statuses

The following extensions to the **wsctx:Status** type MAY be returned by participants and the Context Service to indicate the outcome of executing relevant parts of the protocol; they MAY also be used to indicate the current status of the transaction:

Comment: Place holder as I think Greg has the AI for this under the general issue of asynchronous transactions.

Formatted: Heading 2,H2

Formatted: Normal

Deleted: ¶

Formatted: English (U.K.)

Formatted: Heading 2,H2

Formatted: Bullets and Numbering

Deleted: WS-TXM

Deleted: layers

Deleted: on

Formatted: Heading 2,H2

Formatted: Bullets and Numbering

Comment: This needs reworking because of WS-Context changes.

Formatted: Code

Deleted: *13*

Formatted: Heading 2,H2

Formatted: Bullets and Numbering

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

Deleted: **WS-CTX**

Formatted: Font: Bold

Deleted: are

Deleted: and are also

- RollbackOnly: the status of the endpoint is that it will roll back eventually.
- RollingBack: the endpoint is in the process of rolling back.
- RolledBack: the endpoint has rolled back.
- Committing: the endpoint is in the process of committing. This does not mean that the final outcome will be Committed.
- Committed: the endpoint has committed.
- HeuristicRollback: all of the participants rolled back when they were asked to commit.
- HeuristicCommit: all of the participants committed when they were asked to rollback.
- HeuristicHazard: some of the participants rolled back, some committed and the outcome of others is indeterminate.
- HeuristicMixed: some of the participants rolled back whereas the remainder committed.
- Preparing: the endpoint is preparing.
- Prepared: the endpoint has prepared.

These are specified in the schema, as per Figure 11.

```
<xs:simpleType name="StatusType">
  <xs:restriction base="wstxm:StatusType">
    <xs:enumeration value="activity.status.tx-acid.ROLLBACK_ONLY"/>
    <xs:enumeration value="activity.status.tx-acid.ROLLING_BACK"/>
    <xs:enumeration value="activity.status.tx-acid.ROLLED_BACK"/>
    <xs:enumeration value="activity.status.tx-acid.COMMITTING"/>
    <xs:enumeration value="activity.status.tx-acid.COMMITTED"/>
    <xs:enumeration value="activity.status.tx-
acid.HEURISTIC_ROLLBACK"/>
    <xs:enumeration value="activity.status.tx-acid.HEURISTIC_COMMIT"/>
    <xs:enumeration value="activity.status.tx-acid.HEURISTIC_HAZARD"/>
    <xs:enumeration value="activity.status.tx-acid.HEURISTIC_MIXED"/>
    <xs:enumeration value="activity.status.tx-acid.PREPARING"/>
    <xs:enumeration value="activity.status.tx-acid.PREPARED"/>
  </xs:restriction>
</xs:simpleType>
```

Figure 11, StatusType.

3.8 WS-ACID Faults

This section defines well-known error codes to be used in conjunction with an underlying fault handling mechanism.

Heuristic Hazard

This fault is sent by a participant or the ContextService in response to **wsctx:complete** to indicate that the participant/transaction has terminated in a non-atomic manner. The termination status of all participants (committed or rolled back) is not known.

The qualified name of the fault code is:

```
wsacid:HeuristicHazard
```

Heuristic Mixed

This fault is sent by a participant or the ContextService in response to **wsctx:complete** to indicate that the participant/transaction has terminated in a non-atomic manner. The termination status of some participants was to commit whereas others rolled back.

The qualified name of the fault code is:

Deleted: coordinator or participant

Deleted: coordinator or participant

Deleted: coordinator/participant

Deleted: This may be a transient and in fact, because the protocol uses a presumed-abort optimisation, the NoActivity status can be used to infer that the coordinator cancelled.

Deleted: coordinator/participant

Deleted: coordinator/participant

Deleted: confirmed

Deleted: coordinator/participant

Deleted: coordinator/participant

Deleted: AT

Deleted: Figure 14

Formatted: Code, Don't keep with next

Deleted: *14*

Deleted: *AT*

Comment: Add the structure information for data within wsctx:complete and wsctx:completed messages.

Formatted: Heading 2,H2

Formatted: Font: Bold

```
wsacid:HeuristicMixed
```

Heuristic Rollback

This fault is sent by a participant or the ContextService in response to **wsctx:complete** to indicate that the participant/transaction has rolled back. If this is from the transaction, then all of the participants autonomously rolled back.

The qualified name of the fault code is:

```
wsacid:HeuristicRollback
```

Heuristic Commit

This fault is sent by a participant or the ContextService in response to **wsctx:complete** to indicate that the participant/transaction has committed. If this is from the transaction, then all of the participants autonomously committed.

The qualified name of the fault code is:

```
wsacid:HeuristicCommit
```

Formatted: Heading 2,H2

3.9 Message exchanges

The WS-CAF protocol family is defined in WSDL, with associated schemas. All the WSDL has a common pattern of defining paired port-types, such that one port-type is effectively the requestor, the other the responder for some set of request-response operations.

portType for an initiator ("client" for the operation pair) will expose the responses of the "request/response" as input operations (and should expose the requests as output messages); the responder (service-side) only exposes the request operations as input operations (and should expose the responses as output messages).

Each "response" is shown on the same line as the "request" that invokes it. Where there are a number of responses to a "request", these are shown on successive lines. The initiator portTypes typically include various fault and error operations.

Formatted Table

Initiator (and receiver of response)	**Responder**	**"requests"**	**responses**
CoordinatorParticipant	**Synchronization ParticipantService**	beforeCompletion	beforeCompleted wsctx:InvalidState wsctx:InvalidContext wsctx:NoPermission wsctx:NoContext
		afterCompletion	afterCompleted wsctx:InvalidState wsctx:InvalidContext wsctx:NoPermission wsctx:NoContext

Initiator (and receiver of response)	Responder	"requests"	responses
CoordinatorParticipant	TwoPhase ParticipantService	prepare	voteReadonly voteCommit voteRollback HeuristicMixed HeuristicHazard wsctx:InvalidState wsctx:InvalidContext wsctx:NoContext
		commit	committed HeuristicRollback HeuristicMixed HeuristicHazard wsctx:InvalidState wsctx:InvalidContext wsctx:NoPermission wsctx:NoContext
		rollback	rolledback HeuristicCommit HeuristicMixed HeuristicHazard wsctx:InvalidState wsctx:InvalidContext wsctx:NoPermission wsctx:NoContext
		commitOnePhase	committed rolledback HeuristicMixed HeuristicHazard wsctx:InvalidState wsctx:InvalidContext wsctx:NoPermission wsctx:NoContext
		forgetHeuristic	heuristicForgotten wsctx:InvalidState wsctx:InvalidContext wsctx:NoPermission wsctx:NoContext
		getStatus	status wsctx:InvalidState wsctx:InvalidContext wsctx:NoPermission wsctx:NoContext
wsctx:UserContextService	wsctx:ContextService	wsctx:complete (wsacid:Commit) (wsacid:Rollback)	wsctx:completed (wsacid:Committed) (wsacid:RolledBack) wsacid:HeuristicMixed wsacid:HeurisicHazard

Formatted Table

Formatted: Normal

589

4 References

[1] WSDL 1.1 Specification, see http://www.w3.org/TR/wsdl

[2] "Key words for use in RFCs to Indicate Requirement Levels," RFC 2119, S. Bradner, Harvard University, March 1997.

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

[4] WS-Message Delivery Version 1.0, http://www.w3.org/Submission/2004/SUBM-ws-messagedelivery-20040426/

[5] WS-Reliability latest specification, http://www.oasis-open.org/committees/download.php/8909/WS-Reliability-2004-08-23.pdf. See Section 4.2.3.2 (and its subsection), 4.3.1 (and its subsections). Please note that WS-R defines BareURI as the default.

[6] Addressing wrapper schema, http://www.oasis-open.org/apps/org/workgroup/wsrm/download.php/8365/reference-1.1.xsd

[7] WS-R schema that uses the serviceRefType, http://www.oasis-open.org/apps/org/workgroup/wsrm/download.php/8477/ws-reliability-1.1.xsd

[8] Web Services Addressing, see http://www.w3.org/Submission/ws-addressing/

[9] Web Services Security: SOAP Message Security V1.0, http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-soap-message-security-1.0.pdf

[10] J. N. Gray, “The transaction concept: virtues and limitations”, Proceedings of the 7th VLDB Conference, September 1981, pp. 144-154.

Formatted: Bullets and Numbering

Deleted: [1] WSDL 1.1 Specification, see http://www.w3.org/TR/wsdl¶

Deleted: 2

Deleted:

Formatted: Normal

Formatted: Bullets and Numbering

Appendix A. Acknowledgements

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

Formatted: Bullets and Numbering

Appendix B. 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 implementors or users of this specification, can be obtained from the OASIS Executive Director.

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 Executive Director.

The concepts of atomic transactions have played a cornerstone role in creating today's enterprise application environments by providing guaranteed consistent outcome in complex multiparty business operations and a useful separation of concerns in applications. While numerous multiparty business applications involve various patterns based on atomic transactions in order to solve non-trivial business problems, it was not until recently the word "business transactions" accumulated any concrete meaning. Rapid developments in Internet infrastructure and protocols have yielded a new type of application interoperation concept that makes concepts which could only previously be considered in an abstract form an implementation reality. The effects of such changes have been felt most strongly in business environments, fuelling the mindset for a transition from traditional atomic transactions to extended transaction models better suited for Internet interoperation.

Most business-to-business applications require transactional support in order to guarantee consistent outcome and correct execution. These applications often involve long running computations, loosely coupled systems and components that do not share data, location, or administration and it is thus difficult to incorporate traditional ACID transactions within such architectures. For example, an airline reservation system may reserve a seat on a flight for an individual for a specific period of time, but if the individual does not confirm the seat within that period it will be unreserved. The structuring mechanisms available within traditional transaction systems are sequential and concurrent composition of transactions. These mechanisms are sufficient if an application function can be represented as a single top-level transaction. Frequently with Web services this is not the case. Top-level transactions are most suitably viewed as "short-lived" entities, performing stable state changes to the system; they are less well suited for structuring "long-lived" application functions (e.g., running for minutes, hours, days, …). Long-lived top-level transactions implemented using traditional systems may reduce the concurrency in the system to an unacceptable level by holding on to locks for a long time; further, if such a transaction rolls back, much valuable work already performed could be undone. Web services, because of their inherently unpredictable invocation patterns do not fit well with traditional ACID systems.

2.1Problem statement

As Web Services have evolved as a means to integrate processes and applications at an both inside and outside the firewall, and as Web technologies have become firmly established and widely adopted, traditional transaction semantics and protocols have proven to be inappropriate for some Web services-based applications and services. These particular Web services-based transactions differ from traditional transactions in that they execute over long periods, they require commitments to the transaction to be "negotiated" at runtime, and isolation levels have to be relaxed.

Structuring certain activities from long-running transactions can reduce the amount of concurrency within an application or (in the event of failures) require work to be performed again. For example, there are certain classes of application where it is known that resources acquired within a transaction can be released "early", rather than having to wait until the transaction terminates; in the event of the transaction rolling back, however, certain compensation activities may be necessary to restore the system to a consistent state. Such compensation activities (which may perform

forward or backward recovery) will typically be application specific, may not be necessary at all, or may be more efficiently dealt with by the application.
The goals of the specification are to:
- Provide a basic definition of a core infrastructure service consisting of a Transaction Service for the Web Service environment. The WS-TXM builds on the Web Services Coordination Framework.
- Define the mappings onto the Web Service environment (SOAP message and header definitions, context definition, endpoint address requirements, etc.).
- Define the required infrastructure support such as event mechanisms, etc.
- Define the roles and responsibilities of WS-TXM subcomponents.

3Architecture

Page 6: [2] Deleted **Mark Little** **02/07/2005 10:54 PM**

How application services are invoked is outside the scope of this specification; however, context information related to the sender's activity needs to be referenced and/or propagated.
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.
In order to support both synchronous and asynchronous interactions, the components are described in terms of the behavior and the interactions that occur between them. All interactions are described in terms of correlated messages, which a referencing specification MAY abstract at a higher level into request/response pairs.
Faults and errors that may occur when a service is invoked are communicated back to other Web services in the activity via SOAP messages that are part of the standard protocol. The fault mechanism of the underlying SOAP-based transport isn't used. For example, if an operation fails because no activity is present when one is required, then it will be valid for the InvalidContextFault message to be received by the response service. To accommodate other errors or faults, all response service signatures have a generalFault operation as well as a transientFault operation.
Note, a transientFault message is produced when the implementation finds it cannot successfully execute the requested operation at that time from some *temporary* reason. This reason may be implementation or referencing specification specific. A receiver of a transientFault is free to retry the operation which originally generated it on the assumption that eventually a different response will be produced. Sub-types of transientFault MAY be further defined using the fault model described which can allow for the communication of more specific information on the type of fault.
As long as implementations ensure that the on-the-wire message formats are compliant with those defined in this specification, how the end-points are implemented and how they expose the various operations (e.g., via WSDL [1]) is not mandated by this specification. However, a normative WSDL binding is provided by default in this specification.
Note, this specification does not assume that a reliable message delivery mechanism has to be used for message interactions. As such, it MAY be implementation dependant as to what action is taken if a message is not delivered or no response is received.

Page 9: [3] Deleted **Mark Little** **08/07/2005 11:33 AM**

The activity completion status is either *Failure* or *Success* respectively.

Page 9: [4] Deleted **Mark Little** **05/07/2005 2:17 PM**

; the activity completion status is *Success*

Page 9: [5] Deleted **Mark Little** **05/07/2005 2:17 PM**

, with the activity in the *Failure* completion status

Page 10: [6] Deleted **Mark Little** **05/07/2005 2:18 PM**

In terms of the underlying activity service and coordination service, Figure 6 shows the flow of messages:

1)The application starts a new activity, which ultimately causes the issuing of a *begin* on the Transaction/Coordination Service ALS to demarcate the beginning of a transaction. This causes the Transaction/Coordination Service to create a coordinator used to identify the activity instance and subsequently track the elements interested in the transaction outcome (i.e., the participants).

2)The application issues a server method. Context information is appended to the message. The context is used at the target to recreate the execution environment. Having been passed a coordinator reference the target registers interest in the transaction outcome. (Note, an implementation can choose to register participants directly to the coordinator or through a subordinate coordinator that resides on the target.)

3)The application terminates the activity, which causes a *completeWithStatus* to be sent to the Coordination Service to indicate the end of the transaction. The completion indication is passed to the coordinator. The coordinator sends the two-phase commit protocol messages to each registered participant and returns the outcome to the activity service.

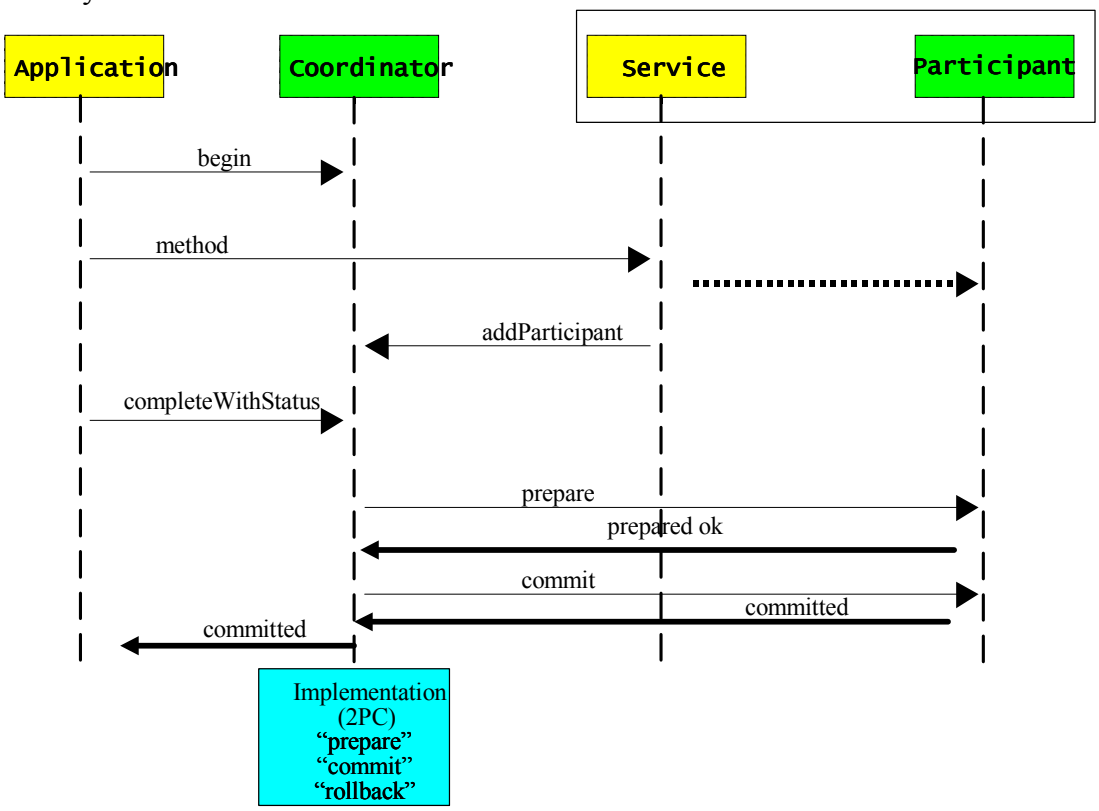

Figure 6, Two-phase commit transactions.

4.3.2Two-phase participant state transitions

The participant state transitions are the same as the coordinators:

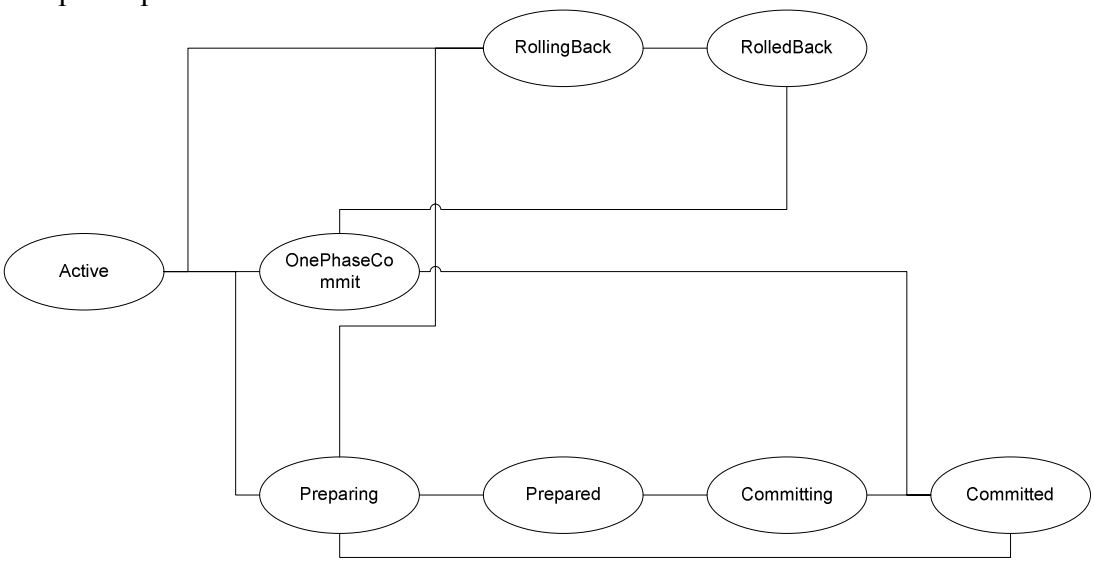

Figure 7, Two-phase participant state transitions.

Two-phase commit

Page 10: [7] Deleted **Mark Little** **05/07/2005 11:26 AM**

because WS-CF supports an asynchronous model

Page 10: [8] Deleted **Mark Little** **05/07/2005 2:27 PM**

. The CoordinatorParticipant end-point address as defined in WS-CF is propagated on all messages

Page 11: [9] Formatted **Mark Little** **08/07/2005 1:43 PM**

Font: Bold

Page 11: [9] Formatted **Mark Little** **05/07/2005 2:31 PM**

Font: Bold

Page 11: [9] Formatted **Mark Little** **05/07/2005 2:31 PM**

Font: Bold

Page 11: [10] Deleted **Mark Little** **05/07/2005 2:28 PM**

it

Page 11: [10] Deleted **Mark Little** **05/07/2005 2:28 PM**

must return

Page 11: [11] Deleted **Mark Little** **08/07/2005 12:19 PM**

Fault

Page 11: [11] Deleted **Mark Little** **05/07/2005 2:28 PM**

message

Page 11: [12] Deleted **Mark Little** **08/07/2005 12:19 PM**

Fault

Page 11: [12] Deleted **Mark Little** **05/07/2005 2:30 PM**

message

Page 11: [13] Deleted **Mark Little** **08/07/2005 12:19 PM**

Fault

Page 11: [13] Deleted **Mark Little** **05/07/2005 2:30 PM**

message

Page 11: [14] Formatted **Mark Little** **05/07/2005 2:30 PM**

Font: Bold

Page 11: [14] Formatted **Mark Little** **05/07/2005 2:30 PM**

Font: Bold

Page 11: [14] Formatted **Mark Little** **05/07/2005 2:31 PM**

Font: Bold

Page 11: [14] Formatted **Mark Little** **05/07/2005 2:32 PM**

Font: Bold

Page 11: [15] Deleted **Mark Little** **05/07/2005 11:32 AM**

confirming

Page 11: [15] Deleted **Mark Little** **05/07/2005 2:32 PM**

If the participant is a subordinate coordinator and cannot determine how all of its enlisted participants terminated then it must return the *heuristicHazardFault* message. If the participant is a subordinate coordinator and some of its enlisted participants rolled back then it must return the *heuristicMixedFault* message.

Page 11: [16] Formatted **Mark Little** **05/07/2005 2:32 PM**

Font: Bold, Not Italic

Page 11: [17] Deleted **Mark Little** **08/07/2005 12:19 PM**

Fault

Page 11: [17] Deleted **Mark Little** **05/07/2005 2:32 PM**

message

Page 11: [18] Deleted **Mark Little** **05/07/2005 2:32 PM**

it is possible for

Page 11: [18] Deleted **Mark Little** **05/07/2005 2:32 PM**

to

Page 11: [18] Deleted **Mark Little** **05/07/2005 2:33 PM**

and not have to execute two phases

Page 11: [18] Deleted **Mark Little** **05/07/2005 2:33 PM**

If the participant is a subordinate coordinator and cannot determine how all of its enlisted participants terminated then it must return the *HeuristicHazardFault* message. If the participant is a subordinate coordinator and some of its

enlisted participants rolled back then it must return the *HeuristicMixedFault* message. If the participant rolls back rather than commits then it must return the **HeuristicRollbackFault** message.

Page 11: [19] Formatted **Mark Little** **05/07/2005 2:34 PM**

Font: Bold, Not Italic