
The Session Concept and Web Services

Hal Hildebrand

Anish Karmarkar

Mark Little

Greg Pavlik

Copyright © 2005 Arjuna Technologies Limited and Oracle Corporation.

November 2005

Abstract

This paper describes the session concept as it relates to middleware systems in general and Web services in particular. Common applications of the session concept are found in distributed object systems, the Web, and messaging middleware systems. The purpose of a session is to allow multiple individual Web Services to enter a relationship by sharing certain common attributes as an externally modeled entity. For example, multiple Web Services executing within the scope of a single authorized/secure session. In the context of Web services, explicit building blocks for session-oriented protocols and services have been proposed in two specifications, WS-Addressing and WS-Context. The distinguishing characteristic of these two proposals is the degree of coupling they introduce between session participants. In this paper we shall compare and contrast the underlying models these specifications present, as they relate to the session concept in Web services. The aim is to illustrate the advantages and disadvantages of both approaches and summarize best-practices and techniques for supporting a scalable Web services architecture.¹

¹Note, although this paper is not purely research oriented, it does make an important contribution in the area of software practices and experiences for current and future researchers. The authors believe that it is important to ensure that the Web services architecture scales as well as the World Wide Web and as we shall see, the session concept and how it is provided play an integral role in that arena.

Table of Contents

1. Introduction	3
2. Session Models in Distributed Systems	3
2.1. CORBA Sessions	4
2.2. HTTP Sessions	5
2.3. Sessions in Message Oriented Middleware	6
3. Web Services Sessions	6
3.1. Web Service Sessions using WS-Addressing	8
3.2. Web Services Sessions using WS-Context	10
4. Comparison of Web Services Session Models	12
4.1. Loosely Coupling	12
4.2. Exposure of Implementation Details	15
4.3. Compatibility with WSDL	15
5. Conclusions	15
References	16

1. Introduction

Agents in distributed systems communicate by sending messages between software components. Middleware software provides plumbing and programming models for applications in order to:

1. Ease the burden of (re)programming network-level message passing on a per application basis.
2. Provide a conceptual framework to organize application logic.
3. Manage system protocols and their relationship to the execution environment.

As an example, distributed object systems like CORBA [\[CORBA\]](#) or Java RMI [\[Java RMI\]](#) provide a programming model for networked services based on the idea of distribution transparency. When a client application works with a distributed object reference, the application itself has no knowledge as to whether the object is collocated in the same address space and available as a local object or distributed across the network. In the latter case, the use of the networked service requires the middleware system to respond to object invocations by marshalling the request to a protocol specific wire-level message and then to send the message over the network. Similarly, the networked service is supported by middleware that decodes the network message and in turn invokes a method on a local object implementation. These kinds of systems provide a simple programming model that is already familiar to developers who have used object-oriented programming languages. The effect is to make building distributed systems relatively simple, at least by comparison to hand-rolling custom message-over-sockets code on a per application basis. Support for system protocols like transactions is layered on top of the basic distributed object model.

One of the common features of all middleware systems is support for the *session concept*. For the purposes of this paper, a session is a mechanism for correlating multiple messages in order to achieve some application-visible semantic. This is typically done on behalf of a client within a service endpoint. In general, middleware systems decouple session association from specific communication channels to improve robustness. To achieve this, the session model is layered on top of a communication channel that links the client to network-visible application services. Many middleware systems advertise the session model explicitly as a mechanism for client applications to manage stateful conversations or communicate with stateful "resources". In other cases, the session concept is maintained less explicitly to support system services that are provided to applications.

At present there is no standard way of obtaining sessions in Web services and this is a significant gap in the architecture. However, there are two specifications that make proposals in this area, WS-Addressing [\[WS-Addressing\]](#) and WS-Context [\[WS-Context\]](#). Although the EndpointReference construct in WS-Addressing is primarily concerned with identifying/addressing Web services (and hence in many ways is complimentary to WS-Context), it provides a session-like concept as a by-product of the model, whereas WS-Context is primarily concerned with sessions. Building upon these specifications and principles to leverage the session concept is an active area of research and development [\[Need for Shared Context\]](#) [\[Comparison of WS-Context and WS-RF\]](#) [\[Modeling Stateful Resources\]](#) [\[WS CDL\]](#) [\[WS Transaction Models\]](#) [\[Business Process Context\]](#). However, as we shall show in this paper, the WS-Addressing session model is based on experiences from previous closely-coupled middleware platforms such as CORBA and J2EE and does not work well in the loosely-coupled environment of Web services. The WS-Context approach is similar in many ways to the Cookie approach in the traditional Web and as such offers an approach that integrates well with the architecture.

2. Session Models in Distributed Systems

In order to understand how the session concept is an integral part of middleware systems, we shall first review several examples in the following sections. We shall then show how the two Web Services approaches (WS-Addressing and WS-Context) build upon some of these core concepts.

2.1. CORBA Sessions

In CORBA systems, the Interoperable Object Reference (IOR) is the primary embodiment of the session concept. The IOR is an address-space distributable reference to an object. The IOR itself is a CDR encoded byte array that contains information of the following form (expressed in CORBA IDL):

```
module IIOP {
  struct Version {
    octet major;
    octet minor;
  };
  struct ProfileBody_1_ {
    Version iiop_version;
    string host;
    unsigned short port;
    sequence<octet> object_key;
  };
  struct ProfileBody_1_1 {
    Version iiop_version;
    String host;
    unsigned short port;
    sequence<octet> object_key;
    sequence<IOP::TaggedComponent> components
  }
}
```

It's important to note that the IOR is encoded not only with the information required to establish a communication channel with a networked service, but also with an opaque token called an Object Key. The Object Key is the foundational session mechanism in CORBA. It is used to demultiplex requests to a specific processing component of the networked service endpoint in CORBA systems. In many cases, the IOR is used to associate requests with individual objects instances that may maintain conversational state on behalf of a client. The Object Key can be supplied by application logic within the service or auto-generated by the infrastructure, and is composed of an arbitrary sequence of bytes: in either case, the IOR originates with the network service.

The structure of the IOR itself is based explicitly on the design of the Object Request Broker infrastructure and reflects the intended implementation model for the system. While often used in support of distributed object systems, the IOR is an example of a *service reference session* model, where the network communication channel and the session identifier are combined into a single data structure that must be dereferenced by clients in combination to access a network service.

A typical application implementation pattern that exploits these features relies on a factory to generate IORs and clients to "garbage collect" resources when they are no longer needed; services may also implement autonomous cleanup in response to inactivity as well. These policies are developed on an ad hoc basis because session expiry is not a part of the basic model. A representative example of this pattern is the stateful session Enterprise JavaBeans (EJB) model [EJB 2.1], where references to stateful session bean instances are used by clients to maintain a stateful conversation with an EJB component. It's worth noting that while EJBs have proven to be highly popular as a mechanism for controlling transactions in J2EE based systems, the stateful session bean model has not been widely deployed. Instead, basic Cookie-based Web sessions have dominated J2EE deployments. Web sessions and stateful EJB sessions are difficult to combine because of their independent session models.

Both EJB and CORBA are best used for tightly coupled systems and assume intimate knowledge of the middleware infrastructure models. Note that in both cases, each objects or components have unique references that must be maintained

by clients for continued use of the networked service. Each reference contains an identifier that is used to dispatch requests to implementation artifacts. This is the defining characteristic of systems that build on service reference sessions.

For this reason, the maintenance of distributed object references is typically complicated. Because sessions and endpoint information is intertwined, broken references must be reconstituted in order to reconnect to existing sessions. In fact, commercial CORBA systems never fully developed an interoperable mechanism that was adopted in practice to provide fault tolerant systems with recoverable session association.

Distributed object systems often support an additional session model, where information that influences the execution semantic of a service is communicated via context information derived from the client execution environment. Distributed transaction processing facilities like the CORBA Object Transaction Service [CORBA] rely on this model to maintain ACID properties across a number of service invocations that involve state changes to shared resources. The generalized view of this execution context-derived session is often referred to as an activity model. The activity model is equivalent in computational expressiveness to the session model used in distributed object references. However, activity sessions have several distinguishing characteristics that make them different than service reference sessions. Activity sessions are:

1. Dynamically bound to the communication channel.
2. Activity sessions may be initiated by clients, by third parties, or by services themselves.
3. Activity sessions incorporate a lifecycle model that assumes a temporal limit. Typically, APIs to support activity sessions assume a timeout exists and often allow clients to manage or specify timeouts.

2.2. HTTP Sessions

Distributed object systems are only one class of middleware frameworks. The ubiquitous HTTP protocol [HTTP 1.1] may be extended to provide sessions for managing application state on behalf of a user agent by the Cookie model [HTTP State Mgmt]. HTTP itself is defined to be a "generic, stateless" protocol capable of acting as a carrier for many domains. This has been particularly useful for supporting information exchange in a massively scalable system like the Web. However, a purely "stateless" protocol has limited applicability to certain kinds of applications. The Cookie model was introduced to allow Web-content applications to flexibly manage state. When a user accesses a Web site's origin server, the session is communicated back to the user agent via a Cookie, which contains information necessary to reestablish the session state on a per-request basis. The content of the Cookie is completely dependent on the origin server. In addition, the origin server is responsible for indicating lifecycle semantics for the Cookie: for example, it may communicate a max-age value for the Cookie or explicitly timeout the Cookie by supplying a max-age value equal to zero on a response message to the user agent. Subsequent access to the Web resources with which the Cookie is linked require the user agent to piggyback the Cookie on the HTTP request headers based on domain name and path information associated with the HTTP request in order to maintain the session semantic.

When application server middleware is used to provide dynamic Web content/applications, middleware infrastructure usually provides built in HTTP Cookie based state management: application developers may store live data in a session cache and the infrastructure handles request correlation with the session data. In this case, the origin server typically ensures that the sessions are bounded by an inactivity time limit. This model is also familiar to developers that are accustomed to working with the Web programming models, for example, the Java Servlet APIs [Java Servlets].

In either case, the HTTP session has no predefined structural relationship with the origin server's Web address or resources. The URL to a web resource represents the resource qua resource and is independent of the Cookie. The server always controls the session and its lifecycle. This model sharply reduces the coupling between user agents and web servers. For example, a Web browser need have no a priori knowledge of Amazon.com in order to purchase a book or CD. Furthermore, each Cookie is tied to a specific Web server, making the construction of a single session that encompasses several different services impossible using this technique.

In general, the flexibility in the HTTP session model exploits the fact that user agents are most often acting on behalf of human users. Recovery from session invalidation is an ad hoc process. While HTTP sessions have been enormously successful, it is not clear that this model may transfer directly to domains concerned with automated machine-to-machine communications.

2.3. Sessions in Message Oriented Middleware

The session model is present in other kinds of middleware systems in less obvious ways. For example, initially it may appear as though message-oriented middleware (MOM) systems only use channels to relay messages to queues or consumers -- that any correlation semantic to backend state must be encoded in the message itself by applications. However, MOM systems offer message-grouping facilities can be applied to ordering and delivery assurance semantics. For example, in many MOM systems, a session is created to demarcate and manage the start of an ordered group. To end delivery of ordered messages within a group, the session is closed. As an example, the Web services standard WS-Reliability [WS-Reliability 1.1] provides explicit protocol instructions to support this paradigm. Sessions are demarcated during application-level message exchange by implicit headers. Message acknowledgements may be communicated on independent communication channels, but messages are correlated with a group identifier that correlates messages with the session. While sessions in MOM systems are not necessarily explicitly accessed by message producers or consumers, the session concept is still very powerful and useful -- including in systems that emphasize nominally decoupled message producers and consumers.

Now that we have described how sessions are defined in other middleware systems, in the next section we shall examine how the session concept is being proposed for Web Services.

3. Web Services Sessions

Before describing how sessions are being added to Web Services, it is important to understand that when we talk about Web services we are making the narrowing assumption that we are talking about services that are described in WSDL [WSDL 1.1] and accessed via SOAP [SOAP 1.1] [SOAP 1.2 Part 1]. We can further make the following observations about the intended design paradigm for the Web services model, which will become important when discussing the alternative approaches to sessions:

1. Web services emphasize loosely coupled systems.
2. Web services do not expose artifacts of the implementation details or component model used to implement the service; Web services emphasize self-encapsulated messages and data transparency.
3. The WSDL definition for a service endpoint is agnostic about session models utilized by the service: specifically, the service element is exposed identically to all clients.

Web Services and state have been a contentious issue for a long time. There are two camps: those that believe Web Services are truly stateless and those that believe they are (or can be) stateful. The advocates against stateful Web Services [REST Web Services] [WS not DS] believe that the presence of state in the server reduces the scalability of the architecture. The argument goes something like this "A *stateful Web Service contains state that is required to conduct a specific unit of work (e.g., reserving a hotel room); however, if it has state then that obviously requires all subsequent users of that state to return to the same Web Service instance. Therefore, if that service endpoint is unavailable, so is the state, which is a limiting factor on scalability. Using stateless Web Services, any state that is required to perform the unit of work must accompany the message that is sent to that service. Therefore, as long as there is a single instance of the type of service that can understand (and deal with) the message, the work can be carried out.*"

However, the proponents of stateful Web Services [Need for Shared Context] [Modeling Stateful Resources] counter with: "A *stateful Web Service has access to state, but that access need not be bound to a single service instance and what happens behind a Web Services endpoint to obtain/manipulate the state is an implementation choice. If it is a requirement that multiple endpoints can use the same state, then that can be supported behind the service endpoints.*"

Support for failover is implicit because of multiple endpoints. All that is required from users and the service implementer is some way of referring to the state on which the work should operate. In fact, the case of using purely stateless services where the state is shipped in each message can be a limiting factor on scalability and performance, particularly as the size of the state to be shipped between endpoints increases."

In reality, both of these state manipulation architectures are valid for a range of applications: in some situations one model may be better suited than another, but there is no clear cut argument for one in favor of the other. State is an inherent property of all Web Services; in fact a purely stateless Web Services (one that neither maintains nor receives state) would have very limited use. The distinguishing hallmark of the Web service paradigm is that services minimize the projection of state into the protocol layers: they favor full encapsulation of the information required to correlate messages with resources and state information within the messages that are included in a Message Exchange Pattern (MEP).

In many, perhaps the majority, of today's use cases for Web services technologies, no explicit session model is required. As a general rule, application messages can maintain sufficient information for service operations to perform the work expressed in the WSDL contract. The emphasis in Web services is on self-contained messages. For example, a Purchase Order message would contain data including a customer id and the purchase order number that are necessary to process the order completely.

However, there are cases in which the session concept is necessary to interject into the execution environment; two examples include the reliable and ordered delivery of a set of messages and the requirement to scope the execution of multiple service operation invocations within a single unit of work. For our purposes, the mechanism for achieving a session in the Web services environment is of fundamental interest. Any use of the session concept increases the coupling between the message producer and the message consumer, so how sessions are represented in the system will have serious implications for the Web services model itself.

In its own way, a standardized session model is as fundamental to the evolution of the Web Services architecture as the standardized adoption of the Cookie model was to the World Wide Web. When the Web was first developed, all server interactions were stateless; sessions between clients and servers are maintained only long enough to transfer an HTML page and are dropped immediately afterward. This means that costly resources, such as operating system network connections, are not maintained for long durations. This per request session model is essential for supporting popular Web sites, where hundred or thousands of users may be interacting with the pages/services it provides.

However, if a site required a client to provide logon credentials before retrieving Web pages, then they would have to be provided on each visit to the site. Fairly obviously this "memory-less" interaction pattern did not scale beyond a few such sites and the Cookies was added to the architecture. As we have noted, Cookies can be used to maintain session-like information between visits to a particular Web site. In fact, session oriented Web services often leverage HTTP Cookies to provide ad hoc session capabilities to support storing application state over multiple message exchanges by exploiting the fact that most Web services are accessed via HTTP(s) today. However, Web services are intended to be accessed via many transport protocols, so this solution is not useable in general.

Currently, there are two proposed mechanisms for modeling the session concept in the Web services community, the EndpointReference in WS-Addressing and the activity concept in WS-Context. As with most aspects of standardization, the value is derived from the potential for its features and functions to be provided by Web Services vendors, therefore helping application developers solve state management and protocol requirements more easily. Once adopted and implemented, the functionality will not only be available as part of the platform but also it will be available in a standard way across platforms, allowing Web Services from multiple environments to interoperate more easily, efficiently, and effectively than if the developers had to code all of the equivalent features and functionality themselves in a non-standard way. The fact that both WS-Addressing and WS-Context are undergoing standardization is a net positive, but we maintain that the session model in WS-Context is a superior building block for the architecture of Web services. We next consider the session embodiment contained in both specifications.

3.1. Web Service Sessions using WS-Addressing

WS-Addressing is a specification that was jointly developed by Microsoft, IBM and others to provide a mechanism for referencing Web services developed on their platforms. In August 2004 the specification was published as a Member Submission to the W3C (a W3C Member Submission called WS-MessageDelivery [WS-MessageDelivery 1.0] was also published in April 2004, offering similar functionality). A W3C Working Group [WS-Addressing WG] was formed in October 2004 to standardize a Web services addressing specification and is using the WS-Addressing Member Submission as the starting point for its work.

WS-Addressing defines two extensible constructs: EndpointReference and Message Information Headers. Message Information Headers provide information on the characteristics of a message that are useful for directing messages to network endpoints and for correlating reply messages. EndpointReferences are used to identify and describe service endpoints and instances. For this paper we will discuss only EndpointReferences, as they are most relevant to the session model in WS-Addressing. EndpointReferences logically extend the WSDL description model. An Endpoint reference contains the following information:

1. An address URI that identifies the endpoint
2. Primary WSDL 1.1 portType of the endpoint
3. WSDL 1.1 service QName/port that contains the definition of the endpoint
4. ReferenceProperties that identify the resource/entity addressed by the EndpointReference. ReferenceProperties are created by the issuer of the EndpointReference and are opaque to the consumer of the EndpointReference.
5. ReferenceParameters that are very similar to ReferenceProperties except that ReferenceParameters are not relevant when comparing two EndpointReference, but ReferenceProperties are. From the point of view of the session model ReferenceParameters do not alter the model and will not be discussed independently in this paper.

The relevant XML schema description for the WS-Addressing EndpointReference looks like:

```
<xs:complexType name="EndpointReferenceType">
  <xs:sequence>
    <xs:element name="Address" type="wsa:AttributedURI"/>
    <xs:element name="ReferenceProperties"
      type="wsa:ReferencePropertiesType" minOccurs="0"/>
    <xs:element name="ReferenceParameters"
      type="wsa:ReferenceParametersType" minOccurs="0"/>
    <xs:element name="PortType"
      type="wsa:AttributedQName" minOccurs="0"/>
    <xs:element name="ServiceName"
      type="wsa:ServiceNameType" minOccurs="0"/>
    <xs:any namespace="other" processContents="lax"
      minOccurs="0" maxOccurs="unbounded">
      <xs:annotation>
        <xs:documentation>
          If "Policy" elements from namespace
            "http://schemas.xmlsoap.org/ws/2002/12/policy#policy"
            are used, they must appear first
            (before any extensibility elements).
        </xs:documentation>
      </xs:annotation>
    </xs:any>
  </xs:sequence>
</xs:complexType>
```



```

</xs:sequence>
<xs:anyAttribute namespace="other" processContents="lax" />
</xs:complexType>

```

In order to better understand the schema description, we provide an example EndpointReference for discussion:

```

<wsa:EndpointReference
  xmlns:wsa="http://schemas.xmlsoap.org/ws/2004/08/addressing"
  xmlns:myns="http://example.com/ref-props-params">
  <wsa:Address>http://example.com/myservice</wsa:Address>
  <wsa:ReferenceProperties>
    <myns:SessionID>session$id:0123456789</myns:SessionID>
  </wsa:ReferenceProperties>
  <wsa:ReferenceParameters>
    <myns:CustomerID>Acme Inc.</myns:CustomerID>
  </wsa:ReferenceParameters>
  <wsa:ServiceName PortName="myPort">myns:myservice</wsa:ServiceName>
</wsa:EndpointReference>

```

In the example above, the endpoint address is "http://myurl.com/myservice/", has a ReferenceProperty containing the session identifier, a ReferenceParameter containing the customer identifier and the WSDL 1.1 service QName and port identifying the description of the endpoint.

The EndpointReference itself appears to be very similar to the IOR structure discussed earlier. The EndpointReference contains network information required to reach a service endpoint in the form of a URL. In addition (and like the IOR), the EndpointReference also contains a set of ReferenceProperties/ReferenceParameters, which are opaque tokens embedded in the reference and propagated back to the service during message dispatch.

SOAP message invocations to the service include each ReferenceProperty and ReferenceParameter bound as a separate SOAP header block. A SOAP request sent to the entity identified by the EndpointReference above looks like:

```

<soap:Envelope
  xmlns:wsa="http://schemas.xmlsoap.org/ws/2004/08/addressing"
  xmlns:soap="http://www.w3.org/2003/05/soap-envelope"
  xmlns:myns="http://example.com/ref-props-params">
  <soap:Header>
    ...
    <wsa:To>http://myurl.com/myservice</wsa:To>
    <myns:CustomerID>Acme Inc.</myns:CustomerID>
    <myns:SessionID>session$id:0123456789</myns:SessionID>
    ...
  </soap:Header>
  <soap:Body>
    ...
  </soap:Body>
</soap:Envelope>

```

WS-Addressing is intended to provide a building block for higher level abstractions. In the specification itself, what exactly an EndpointReference including ReferenceProperties/ReferenceParameters refers to is not defined. It could be an entity modeled by the service, or alternatively, a resource managed by the service. The specification also refers to Web service "instances"; since there is no such entity in the vocabulary of Web services, it's possible the authors of

the WS-Addressing specification intend to release another specification that provides an appropriate definition. The WS-Resource Framework as submitted to OASIS [WSRF TC] was a retrograde set of specifications that built on WS-Addressing to use the base EndpointReference as a building block to structure a model that resembles distributed object systems like CORBA. The WS-RF [WS-RF] specifications provided an elaborate framework that attempts to cope with the brittleness problems created by the reference session model. The design intention of WS-RF was to provide a general model for stateful Web services [Modeling Stateful Resources].

Because WS-Addressing is clearly based on the session reference model², a natural way to think about the ReferenceProperties/ReferenceParameters is to compare them to the Object Key in the IOR found in CORBA systems. The ReferenceProperties/ReferenceParameters provide a mechanism for reestablishing the execution context for the request message once it reaches the service network endpoint. Because WS-Addressing is a "building block" specification for protocols and product features, it is positioned to provide a fundamental session model for the Web services architecture moving forward.

3.2. Web Services Sessions using WS-Context

Since December of 2003, the Web services community has been developing a standard called WS-Context in the OASIS consortium [WS-CAF TC] to provide an explicit session model for Web services. WS-Context provides a session construct that is decoupled from application-domain service representations.

The specification defines a basic context structure, explains how it can be linked explicitly to the activity concept, and provides WSDL bindings for services that manage activity-sessions. WS-Context also defines the notion of an activity to which the context is bound: the lifetime of the activity is the lifetime of the context. An activity models a session: all interactions on a session-oriented service in the scope of an activity will be uniquely and unambiguously tied to that activity through the context. An activity may be used to model specific session oriented concepts like a business process instance or conversation.

The specification is decomposable, so that session models that don't map precisely to the activity concept can use the context structure as well; currently proposed specifications that use the activity session model include WS-Reliability [WS-Reliability 1.1] and WS-ReliableMessaging [WS-ReliableMessaging], WS-Coordination [WS-Coordination] and WS-CoordinationFramework [WS-CAF TC], WS-AtomicTransaction [WS-AtomicTransaction], WS-BusinessActivity [WS-BusinessActivity], WS-TXM [WS-CAF TC], SOAPConversations [SOAP Conversation] and WS-Enumeration [WS-Enumeration], suggesting that the model is particularly robust for building session based protocols.

The WS-Context context structure looks like:

```
<xsd:complexType name="ContextType">
  <xsd:sequence>
    <xsd:any namespace="##other" processContents="lax"
      minOccurs="0" maxOccurs="unbounded"/>
    <xsd:element name="context-identifier"
      type="tns:contextIdentifierType"/>
    <xsd:element name="context-service"
      type="ref:ServiceRefType" minOccurs="0"/>
    <xsd:element name="type" type="xsd:anyURI"/>
    <xsd:element name="context-manager"
      type="ref:ServiceRefType" minOccurs="0"/>
    <xsd:element name="parent-context"
      type="tns:ContextType" minOccurs="0">
  </xsd:sequence>
```

²Please note that an EndpointReference can indeed be created and used without ReferenceProperties/ReferenceParameters. Such a usage does not constitute a session reference model. This paper deals only with EndpointReferences that contain ReferenceProperties/ReferenceParameters.

```
<xsd:attribute name="timeout" type="xsd:int" use="optional" />
<xsd:attribute ref="wsu:Id" use="optional" />
</xsd:complexType>
```

There are several important aspects of the Context structure:

1. The context provides a required type attribute to identify the type of session model that is being represented as the intended interaction semantic. This allows processors to unambiguously understand the "meaning" of the contextualization header.
2. Contexts contain explicit lifecycle semantics. A time to live is normally associated with sessions. This provides a normal framework for understanding the validity and invalidation rules for a session.
3. The context can contain arbitrary data elements. Depending on the rules governing the system and protocols in use, these data elements may be arbitrarily augmented by the clients and service endpoints. This allows for sessions to be shared across multiple services and for the shared session content to be change over time.
4. Contexts may be passed by-value or by-reference. By-value contexts support familiar use cases like distributed transaction processing. By-reference contexts are particularly useful for business protocols that require large amounts of data to be included with messages, for example, legal restrictions related to a sales process.
5. The Context structure acts as a wrapper for information related to the session protocol. The SOAP binding for WS-Context includes the Context itself as a SOAP header, maintaining clear encapsulation rules for contextualized messages.

These characteristics allow the Context to be used for a variety of protocols, including the direct management of application state. Most importantly, the session concept is loosely coupled with respect to communication channels and service endpoints: the session may be used in conjunction with a service for a short period or even shared across multiple services. Late binding also means that protocols may use WS-Context to support either ephemeral or long-lived sessions associated with a fixed service endpoint definition as appropriate within an application.

As mentioned, the context is bound to SOAP messages as a header. An example SOAP message including a context header would appear on the wire as:

```
<soap:Envelope xmlns:soap="http://www.w3.org/2002/06/soap-envelope">
  <soap:Header>
    <context xmlns="http://docs.oasis-open.org/wscaf/2004/09/wsctx"
      timeout="100" soap:mustUnderstand="1">
      <context-identifier>
        http://example.com/abcdef:012345
      </context-identifier>
      <type>http://example.com/context/type1</type>
    </context>
  </soap:Header>
  <soap:Body>
    <!-- Application Payload -->
  </soap:Body>
</soap:Envelope>
```

The WS-Context specification also provides two (optional) service endpoint interfaces for managing the lifecycle and contents of by-reference contexts, the ContextService and the ContextManager. The ContextService provides the ability to begin activities and terminate (complete) activities. It's important to note that this interface is optional: services

could create, use and manage Contexts like HTTP Cookies. The ContextManager allows agents to dereference Contexts for up-to-date data content and is used in conjunction with by-reference Contexts.

The target use cases for WS-Context include support for transaction models in the Web services environment, reliable messaging sessions, structured/global interaction models like choreographies and conversational state management in temporally coupled service interactions. It is intended to be used as a building block by other specifications that require session constructs -- in fact, several specifications related to transaction protocols in OASIS are already building on WS-Context -- and to supply execution contexts for application level state correlation. In the latter case, instead of adopting a reference-based session model like CORBA, the application session model can be used in ways that are largely analogous to HTTP sessions that are represented by a Cookie.

WS-Context provides a more generalized session model than what is present in WS-Addressing and is focused on session support with minimal coupling between services. The WS-Context model is also consistent with the model of Web services provided by WSDL: all clients have an equivalent view of the service, references to which are based on information provided in the WSDL service element.

Having examined both approaches to the session model for Web Services, in the following section we shall compare and contrast them. The aim is to attempt to show which of these approaches is most applicable for an environment like the Web, which must be highly scalable and tolerate frequent failures. This is obviously important when considering which model to leverage in future research and development.

4. Comparison of Web Services Session Models

As we have seen, currently there are two primary models for the session concept that are being defined to advance the next generation Web services architecture: the WS-Addressing EndpointReference with ReferenceProperties/ReferenceParameters and the WS-Context explicit context structure. The WS-Addressing session model provides coupling between the Web service endpoint information and the session data, which is analogous to object references in distributed object systems. WS-Context provides a session model that is an evolution of the session models found in HTTP servers, transaction, and MOM systems.

As we observed, the fundamental, widely-accepted characteristics of Web services are that they can be described in WSDL and are accessed via SOAP. WS-Addressing EndpointReferences plus ReferenceProperties/ReferenceParameters do not map to any concepts described in WSDL. WS-Context allows the session requirements to be described in the WSDL SOAP bindings, by specifying the necessary headers. Capabilities related to sessions can also be described in a policy description for the service.

We also considered three other characteristics that made Web Services attractive for integrating computing systems: Loose coupling, exposure of implementation details, and compatibility with WSDL. In the following sections we shall examine how both session models affect these characteristics.

4.1. Loosely Coupling

The Web Services architecture is an inherently loosely coupled system, with the capability (if architected correctly) to offer scalability, flexibility and extensibility on a par with the World Wide Web. The interaction pattern for Web Services is based on coarse-grained services or components. The architecture is deliberately not prescriptive about what happens behind service endpoints: Web Services are ultimately only concerned with the transfer of structured data between parties, plus any meta-level information to safeguard such transfers (e.g., by encrypting or digitally signing messages). This gives flexibility of implementation, allowing systems to adapt to changes in requirements, technology etc. without directly affecting users. It also means that issues such as whether or not a service maintains state on behalf of users or their (temporally bounded) interactions, has been an implementation choice not typically exposed to users.

Unfortunately the WS-Addressing EndpointReferences with ReferenceProperties/ReferenceParameters approach tightly couple sessions to references. Clients *cannot* switch or alter the interaction semantic with respect to the service. Clients *must* maintain a special reference on a per-relationship basis with each service, further coupling the service client and the service itself. This has two important consequences: it creates a brittle relationship between the client and the network service in which the client's understanding of the service is limited to a particular session. Termination of that session invalidates the client's communication channel to the service. Secondly, this results in a scalability problem: clients must contain special reference-pointers to services for each relationship that is linked by the session. Often this results in the unnecessary management and storage of redundant data.

As in other distribution environments where the reference session model is used (e.g., CORBA or J2EE), the remote reference (address) that the client has to the service endpoint *must* be remembered by the client for subsequent invocations. If the client application interacts with multiple services within the same logical session, then it is often the case that the state of a service has relevance to the client only when used in conjunction with the associated states of the other services. This necessarily means that the client must remember each service reference and somehow associate them with a specific interaction; multiple interactions will obviously result in different reference sets that may be combined to represent each sessions.

For example, if there are N services used within the same application session, each maintaining m different states, the client application will have to maintain $N*m$ reference endpoints. It is worth remembering that the initial service endpoint references will often be obtained from some bootstrap process such as UDDI. But in this model, these references are stateless and of no use beyond starting the application interactions. Subsequent visits to these sites that require access to specific states must use different references in the WS-Addressing model.

Consider the example shown in [Figure 1, "Context and scalable session management."](#) There are three Web Services that offer services that implicitly manipulate state on behalf of users. In one execution of the application, the user's interactions create and manipulate states A, B and D, whereas in another execution states C and E are used.

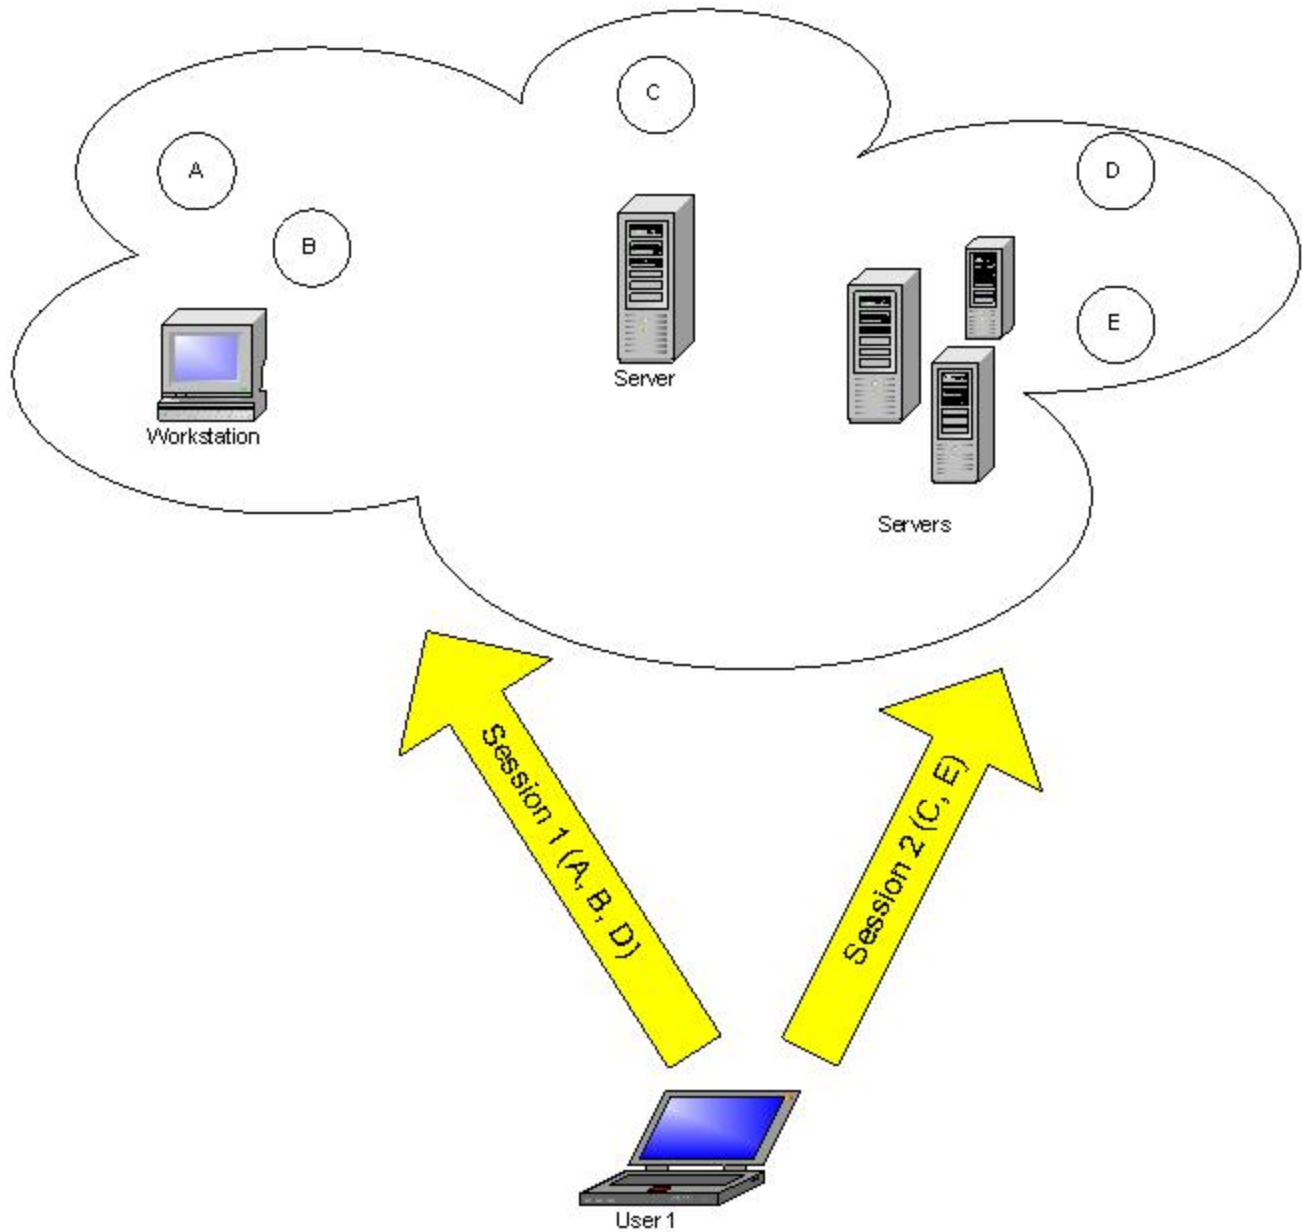


Figure 1. Context and scalable session management.

Maintaining each EndpointReference for every session in order to obtain the same states on subsequent invocations obviously does not scale to an environment the size of the Web.

On the other hand, WS-Context allows a service client to more naturally bind the relationship to the service dynamically and temporarily. The client's communication channel to the service is not impacted by a specific session relationship. This has special implications as we consider scaling Web services from intra-domain deployments to general services offered on the Internet.

The WS-Context approach continues to embrace the inherently loosely-coupled nature of Web Services. As we have shown, each interaction with a set of services can be modeled as a session, and this in turn can be modeled as a WS-Context activity with an associated context. Using context offers scalability in that no matter how many states may be

used by a complex application, they are implicitly referred to by a single context; the user need only remember the context (session) identifier in order to be able to use them in subsequent invocations.

Whenever a client application interacts with a set of services within the same session, the context is propagated to the services and they map this context to the necessary states that the client interaction requires. How this mapping occurs is an implementation specific choice that need not be exposed to the client. Changes to this implementation choice (e.g., load balancing the states across different servers) does not affect users. Furthermore, since each service within a specific session gets the same context, upon later revisiting these services and providing the same context again, the client application can be sure to return to a consistent set of states. So for the N services and m states in our previous example, the client need only maintain N endpoint references and as we mentioned earlier, typically these will be obtained from the bootstrap process anyway. Thus, this model offers much better scalability.

4.2. Exposure of Implementation Details

Web services do not expose artifacts of the implementation details or component model used to implement the service. WS-Addressing `EndpointReferences` with `ReferenceProperties/ReferenceParameters` naturally encourage modeling components and/or resources behind the service façade by focusing attention on pointers to something other than the service itself.

The utility of this design is the ability to export references to abstractions with memory across invocations, rather than services themselves. This shifts the design center from self-contained messages and toward operations that offer functionality based on the identity encoded in the session reference. Consider a purchase order. A message-centric system might accept a purchase order as a self-contained document. In this case, WSDL operations would require purchase orders as input messages for a few coarse grained functions. Using the session reference model, a purchase order might be exposed to users as an `EndpointReference` to a purchase order managed by the system. Operations would tend to be fine-grained modifications to the properties of the purchase order identified by the reference that induce state changes on an internal representation of a purchase order abstraction.

WS-Context focuses on the expression of a session semantic only for the duration of the active session. It bears no relationship to the backend resources or implementation strategy employed by the service. Even if used to manage conversational state explicitly, there is no built-in relationship between WS-Context and any domain-specific modeling artifacts: the session mechanism and domain models are considered to be orthogonal.

4.3. Compatibility with WSDL

WSDL itself is agnostic about session models: nothing about the semantic of the session model contained in `EndpointReferences` plus `ReferenceProperties/ReferenceParameters` can be expressed in WSDL. At best, a factory pattern can be used to return object-reference-like structures to clients. From this perspective, WS-Addressing is broadly misaligned with the normal models provided for developing Web services.

Requirements for WS-Context headers can be expressed directly in WSDL and map naturally to other proposed mechanisms for expressions about the capabilities and requirements of services like WS-Policy.

5. Conclusions

Design decisions for service implementers are strongly influenced by how session concepts are expressed in the Web services environment. Of the currently proposed models available to capture the session concept, WS-Context focuses on loosely coupled, message-oriented systems design. We believe this is the appropriate model to ensure that Web services successfully fulfill their role as a technology for integrating heterogeneous computing environments with minimal coupling between systems. Furthermore we believe that context and its use in state management is fundamental to maintaining the inherently loosely coupled nature of the Web Services architecture. It provides a scalable way of referring to state without breaking the Web Services component abstraction: what happens within a Web Service is

purely implementation and does not need to be exposed to users. To do otherwise can reduce scalability and lead to tightly coupled applications.

References

- [SOAP 1.1] *Simple Object Access Protocol (SOAP) 1.1*, D. Box et al, W3C Note, 08 May 2000. Available at <http://www.w3.org/TR/2000/NOTE-SOAP-20000508/> .
- [WS-Addressing] *Web Services Addressing (WS-Addressing)*, D. Box et al, W3C Member Submission, 10 August 2004. Available at <http://www.w3.org/Submission/2004/SUBM-ws-addressing-20040810/> .
- [WSDL 1.1] *Web Services Description Language (WSDL) 1.1*, E. Christensen et al, W3C Note, 15 March 2001. Available at <http://www.w3.org/TR/2000/NOTE-SOAP-20000508/> .
- [WS-RF] *The WS-Resource Framework Version 1.0*, K. Czajkowski et al, 05 March 2004. Available at <http://www-fp.globus.org/wsrf/specs/ws-wsrf.pdf> .
- [EJB 2.1] *Enterprise JavaBeans Specification, Version 2.1*, L.G. DeMichiel, Specification Lead, 12 November 2003. Available at <http://jcp.org/aboutJava/communityprocess/final/jsr153/index.html> .
- [HTTP 1.1] *Hypertext Transfer Protocol -- HTTP/1.1*, R. Fielding et al, IETF Network Working Group Request for Comments: 2616, June 1999. Available at <http://www.ietf.org/rfc/rfc2616.txt> .
- [Modeling Stateful Resources] *Modeling Stateful Resources with Web Services*, I. Foster, et al, 05 March 2004. Available at <http://www-106.ibm.com/developerworks/library/ws-resource/ws-modelingresources.pdf> .
- [SOAP 1.2 Part 1] *SOAP Version 1.2 Part 1: Messaging Framework*, M. Gudgin et al, W3C Recommendation, 24 June 2003. Available at <http://www.w3.org/TR/2003/REC-soap12-part1-20030624/> .
- [Java Servlets] *Java Servlet Programming* J. Hunter January 2001 Available at <http://www.amazon.com>
- [WS-Reliability 1.1] *WS-Reliability 1.1*, K. Iwasa et al, OASIS Standard, 15 November 2004. Available at http://docs.oasis-open.org/wsrn/ws-reliability/v1.1/wsrn-ws_reliability-1.1-spec-os.pdf .
- [WS-ReliableMessaging] *WS-ReliableMessaging*, D. Langworthy (Editor), February 2005. Available at <http://msdn.microsoft.com/library/en-us/dnglobspec/html/wsrmspecindex.asp>
- [WS-MessageDelivery 1.0] *WS-MessageDelivery Version 1.0*, A. Karmarkar et al, W3C Member Submission, 26 April 2004. Available at <http://www.w3.org/Submission/2004/SUBM-ws-messagedelivery-20040426/> .
- [HTTP State Mgmt] *HTTP State Management Mechanism*, D. Kristol et al, IETF Network Working Group Request for Comments: 2965, October 2000. Available at <http://www.ietf.org/rfc/rfc2965.txt> .
- [WS-Context] *Web Services Context Specification (WS-COContext)*, M. Little et al, OASIS Committee Draft version 0.8, 25 October 2004. Available at <http://www.oasis-open.org/committees/download.php/9806/> .
- [WS-CAF TC] *OASIS Web Services Composite Application Framework (WS-CAF) TC*. Available at http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=ws-caf .
- [WS-Coordination] *WS-Coordination Specification Version 1.0*, M. Feingold (Editor), August 2005. Available at <http://www.arjuna.com/library/specs/ws-tx/WS-Coordination.pdf>
- [WS-AtomicTransaction] *WS-AtomicTransaction Specification Version 1.0*, M. Feingold (Editor), August 2005. Available at <http://www.arjuna.com/library/specs/ws-tx/WS-AtomicTransaction.pdf>

[WS-BusinessActivity] *WS-BusinessActivity Specification Version 1.0*, M. Feingold (Editor), August 2005. Available at <http://www.arjuna.com/library/specs/ws-tx/WS-BusinessActivity.pdf>

[WSRF TC] *OASIS Web Services Resource Framework (WSRF) TC*. Available at http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=wsrf .

[CORBA] *Common Object Request Broker Architecture: Core Specification*, Object Management Group, March 2004. Available at <http://www.omg.org/docs/formal/04-03-01.pdf> .

[Java RMI] *Java Remove Method Invocation Specification*, Sun Microsystems, 2003. Available at <http://java.sun.com/j2se/1.4.2/docs/guide/rmi/spec/rmiTOC.html> .

[WS-Addressing WG] *W3C Web Services Addressing Working Group*. Available at <http://www.w3.org/2002/ws/addr/> .

[Comparison of WS-Context and WS-RF] *Stateful Interactions in Web Services: a comparison of WS-Context and WS-Resource Framework*, J. Webber et al, Web Services Journal, May 2004.

[Need for Shared Context] *Enabling Open, Interoperable, and Smart Web Services (The Need for Shared Context)*, A. T. Manes, W3C workshop on Web services, 12 March 2001. Available at <http://www.w3.org/2001/03/WSWS-popa/paper29> .

[WS CDL] *Web Services Choreography Description Language Version 1.0*, N. Kavantzias et al, W3C Working Draft, 27 April 2004. Available at <http://www.w3.org/TR/2004/WD-ws-cdl-10-20040427/> .

[WS-Enumeration] *WS-Enumeration Specification*, A. Geller (Editor), September 2004. Available at <http://msdn.microsoft.com/ws/2004/09/ws-enumeration/>

[WS Transaction Models] *Models for Web services transactions*, M. Little, SIGMOD Paris, 2004.

[SOAP Conversation] *SOAP Conversation Protocol Version 1.0*, D. Bau and D. Orchard, June 2002. Available at <http://dev2dev.bea.com/pub/a/2002/06/SOAPConversation.html>

[Business Process Context] *Business Process Context: Context, coordination and transactions in the Web Services architecture*, D. Bunting et al, XML 2004 Conference, November 2004. Available at http://www.idealliance.org/proceedings/xml04/papers/184/XML2004_WS-Caf_paper.html .

[REST Web Services] *Building Web Services the REST Way*, R. L. Costello. Available at <http://www.xfront.com/REST-Web-Services.html> .

[WS not DS] *Web Services Are Not Distributed Objects*, W. Vogels, 26 August 2003. Available at <http://weblogs.cs.cornell.edu/AllThingsDistributed/archives/000120.htm> .

Biography

Hal Hildebrand

[Oracle Corporation](http://www.oracle.com) [http://www.oracle.com]

Moss Beach

California

United States of America

Hal.Hildebrand@oracle.com [mailto:Hal.Hildebrand@oracle.com]

Hal Hildebrand is currently the Architect of Oracle's Java Product Group. He has been producing product from specifications for many years at Oracle.

Anish Karmarkar

[Oracle Corporation](http://www.oracle.com) [http://www.oracle.com]

Portland

Oregon

United States of America

Anish.Karmarkar@oracle.com [mailto:Anish.Karmarkar@oracle.com]

Dr. Anish Karmarkar is a Principal Member of Technical Staff in the Core Web Services and Standards Group at Oracle Corporation. He has been working on Distributed Systems for the last 13 years. He is a member of the WS-Addressing Working Group in W3C. He is also a member of various Web services related Working Groups, Technical Committees, Expert Groups in the W3C, OASIS, WS-I and JCP. He has coauthored various Web services related specifications and test suites.

Mark Little

[Arjuna Technologies Limited](http://www.arjuna.com) [http://www.arjuna.com]

Newcastle upon Tyne

England

United Kingdom

Mark.Little@arjuna.com [mailto:Mark.Little@arjuna.com]

Dr. Mark Little is Chief Architect for Arjuna Technologies Ltd, a company specialising in the development of reliable middleware. Before this, Mark was a Distinguished Engineer/Architect within HP Arjuna Labs. in Newcastle upon Tyne, England, where he lead the HP-TS and HP-WST teams, developing J2EE and Web services transactions products respectively. He is one of the primary authors of the OMG Activity Service specification and is on the expert group for the same work in J2EE (JSR 95). He is also the specification lead for JSR 156: Java API for XML Transactions. He is on the OTS Revision Task Force and the OASIS Business Transactions Protocol specification. Before joining HP he was for over 10 years a member of the Arjuna team within the University of Newcastle upon Tyne (where he continues to have a Visiting Fellowship). His research within the Arjuna team included replication and transactions support, which include the construction of an OTS/JTS compliant transaction processing system. Mark has published extensively in the Web Services Journal, Java Developers Journal and other journals and magazines. He is also the co-author of several books including "Java and Transactions for Systems Professionals" and "The J2EE 1.4 Bible."

Greg Pavlik

[Oracle Corporation](http://www.oracle.com) [http://www.oracle.com]

Mooretown

New Jersey

United States of America

Greg.Pavlik@oracle.com [mailto:Greg.Pavlik@oracle.com]

Greg Pavlik has been working on the design and implementation of Java middleware systems for the last 8 years. He is currently Web Services Architect for Oracle Corporation, a role in which he focuses on a combination of

technology strategy, product development and standards work. Before joining Oracle, Greg was a Distinguished Engineer at Hewlett Packard Corporation and Chief Architect of the Application Server Division. Greg has been involved in a number of Web services standards, including as a member of the Web Services Remote Portlet (WSRP) Technical Committee, the Web Services Notification Technical Committee (WS-N), the Web Services Resource Framework (WS-RF), and the WS-Composite Application Framework (WS-CAF). He is currently one of the editors of both the WS-Context and WS-Coordination Framework specifications.