



Unraveling the Web Services Web

An Introduction to SOAP, WSDL, and UDDI

Francisco Curbera, Matthew Duffler, Rania Khalaf, William Nagy, Nirmal Mukhi, and Sanjiva Weerawarana • IBM T.J. Watson Research Center

Over the past few years, businesses have interacted using ad hoc approaches that take advantage of the basic Internet infrastructure. Now, however, Web services are emerging to provide a systematic and extensible framework for application-to-application interaction, built on top of existing Web protocols and based on open XML standards.

Say, for example, that you want to purchase a vacation package using an online travel agent. To locate the best prices on airline tickets, hotels, and rental cars, the agency will have to poll multiple companies, each of which likely uses different, incompatible applications for pricing and reservations. Web services aim to simplify this process by defining a standardized mechanism to describe, locate, and communicate with online applications. Essentially, each application becomes an accessible Web service component that is described using open standards. An online travel service could thus use the same Web services framework to locate and reserve your package elements, as well as to lease Internet-based credit check and bank payment services on a pay-per-use basis to expedite fund transfers between you, the travel agency, and the vendors.

The Web services framework is divided into three areas – communication protocols, service descriptions, and service discovery – and specifications are being developed for each. In this article, we look at the specifications that are currently the most salient and stable in each area:

- the simple object access protocol (SOAP, www.w3.org/2000/xp) which enables communication among Web services;
- the Web Services Description Language (WSDL, www.w3.org/TR/wsdl.html), which provides a formal, computer-readable description of Web services; and
- the Universal Description, Discovery, and Integration (UDDI, www.uddi.org) directory, which is a registry of Web services descriptions.

At this point, Web services technology is still emerging, and researchers are still developing important pieces, including quality of service descriptions and interaction models. Because the Web services framework is modular, however, you can use just the parts of the stack you need. Therefore, developers can take advantage of the available specifications and tooling now and incorporate more modules as the technology matures.

Communication: SOAP

Given the Web's intrinsically distributed and heterogeneous nature, communication mechanisms must be platform-independent, international, secure, and as lightweight as possible. XML is now firmly established as the lingua franca for information and data encoding for platform independence and internationalization. Building a communication protocol using XML is thus a natural answer for Web services.

Enter SOAP, which was initially created by Microsoft and later developed in collaboration with Developmentor, IBM, Lotus, and UserLand. SOAP is an XML-based protocol for messaging and remote procedure calls (RPCs). Rather than define a new transport protocol, SOAP works on existing transports, such as HTTP, SMTP, and MQSeries.

At its core, a SOAP message has a very simple

structure: an XML element with two child elements, one of which contains the header and the other the body. The header contents and body elements are themselves arbitrary XML. Figure 1 shows a SOAP envelope's structure.

In addition to the basic message structure, the SOAP specification defines a model that dictates how recipients should process SOAP messages. The message model also includes *actors*, which indicate who should process the message. A message can identify actors that indicate a series of intermediaries that process the message parts meant for them and pass on the rest.

Messaging Using SOAP

At the basic functionality level, you can use SOAP as a simple messaging protocol. Throughout this article, we'll illustrate the Web services specifications using a simple example drawn again from the travel services industry. Our traveler, Joe, is scheduled for an afternoon flight and wants to checkin electronically. We'll assume that Joe knows of a service with an electronic `CheckIn` method and that he knows the format for encoding the ticket. Given this, he could simply create and send a SOAP message to that service for processing.

Figure 2 shows such a SOAP message, carried by HTTP. The HTTP headers are above the `SOAP:Envelope` element. The `POST` header shows that the message uses HTTP `POST`, which browsers also use to submit forms. Following the `POST` header is an optional `SOAPAction` header that indicates the message's intended purpose. If there were a response, the HTTP response would be of type `text/xml`, as declared in the `Content-Type` header, and could contain a SOAP message with the response data. Alternatively, the recipient could deliver the response message later (asynchronously).

Note that the message in Figure 2 has no SOAP headers; the body simply contains an XML representation of an e-ticket with the person's name and flight details. Any realistic B2B scenario would, of course, have many headers indicating further information, including the sender's credentials and correlation information.

Remote Procedure Calls Using SOAP

To use SOAP for RPCs, you must define an RPC protocol, including:

- how typed values can be transported back and forth between the SOAP representation (XML) and the application's representation (such as a Java class for a ticket), and

```
<SOAP:Envelope xmlns:SOAP=
    "http://schemas.xmlsoap.org/soap/envelope/">
  <SOAP:Header>
    <!-- content of header goes here -->
  </SOAP:Header>
  <SOAP:Body>
    <!-- content of body goes here -->
  </SOAP:Body>
</SOAP:Envelope>
```

Figure 1. Structure of a SOAP message. The envelope features child elements that contain the message header and body elements.

```
POST /travelservice
SOAPAction: "http://www.acme-travel.com/checkin"
Content-Type: text/xml; charset="utf-8"
Content-Length: nnnn

<SOAP:Envelope xmlns:SOAP=
    "http://schemas.xmlsoap.org/soap/envelope/">
  <SOAP:Body>
    <et:eTicket xmlns:et=
        "http://www.acme-travel.com/eticket/schema">
      <et:passengerName first="Joe" last="Smith"/>
      <et:flightInfo airlineName="AA"
        flightNumber="1111"
        departureDate="2002-01-01"
        departureTime="1905"/>
    </et:eTicket>
  </SOAP:Body>
</SOAP:Envelope>
```

Figure 2. SOAP message containing an e-ticket. The SOAPAction header indicates the message's purpose. In a real-world scenario, the message would contain additional information, including the sender's credentials.

- where the various RPC parts are carried (object identity, operation name, and parameters).

The W3C's XML schema specification (www.w3.org/XML/Schema) provides a standard language for defining the document structure and the XML structures' data types. That is, given a type like `integer` or a complex type, such as a record with two fields (say, an `integer` and a `string`), XML schema offers a standard way to write the type in XML. To enable transmission of the typed values, SOAP assumes a type system based on the one in XML schema and defines its canonical encoding in XML. Using this encoding style, you can produce an XML encoding for any type of structured data. RPC arguments and responses are also rep-

```

POST /travelservice
SOAPAction: "http://www.acme-travel.com/flightinfo"
Content-Type: text/xml; charset="utf-8"
Content-Length: nnnn

<SOAP:Envelope xmlns:SOAP=
  "http://schemas.xmlsoap.org/soap/envelope/">
  <SOAP:Body>
    <m:GetFlightInfo
      xmlns:m="http://www.acme-travel.com/flightinfo"
      SOAP:encodingStyle=
        "http://schemas.xmlsoap.org/soap/encoding/"
      xmlns:xsd="http://www.w3.org/2001/XMLSchema"
      xmlns:xsi=
        "http://www.w3.org/2001/XMLSchema-instance">
      <airlineName xsi:type="xsd:string">UL
    </airlineName>
      <flightNumber xsi:type="xsd:int">506
    </flightNumber>
    </m:GetFlightInfo>
  </SOAP:Body>
</SOAP:Envelope>

```

Figure 3. SOAP RPC call. To find out if his flight is on time, Joe sends a string containing the airline's name and an integer with the flight number.

resented using this encoding.

Now, let's say Joe wants to know whether his flight has been delayed. He knows that the service has a function, `GetFlightInfo`, which takes two arguments – a string containing the airline name and an integer with the flight number – and returns a structured value (a record) with two fields – the gate number and flight status. In this case, Joe can get flight status by sending the service an HTTP POST carrying a SOAP envelope like the one in Figure 3.

In this SOAP envelope, the call to `GetFlightInfo` is an XML element with attributes that include information about the encoding (note the references to XML schema). The element's children are the method call's arguments: `airlineName` and `flightNumber`. Their types are defined in the type attributes, where `xsd` refers to the XML schema definitions. When the SOAP implementation receives the message, it converts the XML text for `UL` and `506` into the appropriate string and integer based on the service's implementation. It then calls the `GetFlightInfo` method with those arguments.

Figure 4 shows the response to this request. In this case, the response contains a structured value with the subvalues for gate number and flight status. Luckily, Joe's flight is on time.

SOAP implementations exist for several programming languages, including C, Java, and Perl,

which automatically generate and process the SOAP messages. Assuming the messages conform to SOAP specifications, they can thus be exchanged by services implemented in different languages.

Description: WSDL

Speaking a universal language is not very useful unless you can maintain the basic conversations that let you achieve your goals. For Web services, SOAP offers basic communication, but it does not tell us what messages must be exchanged to successfully interact with a service. That role is filled by WSDL, an XML format developed by IBM and Microsoft to describe Web services as collections of communication end points that can exchange certain messages. In other words, a WSDL document describes a Web service's interface and provides users with a point of contact.

In this section, our examples are fragments of a WSDL document that describes a Web service that can process the two types of interactions in our SOAP examples. The first interaction, `GetFlightInfo`, is accessed using the SOAP RPC model; it takes an airline name and a flight number and returns a complex (or structured) type with flight information. The second, `CheckIn`, uses pure SOAP messaging; it expects to receive an XML representation of an electronic ticket, and returns no information.

A complete WSDL service description provides two pieces of information: an application-level service description, or abstract interface, and the specific protocol-dependent details that users must follow to access the service at concrete service end points. This separation accounts for the fact that similar application-level service functionality is often deployed at different end points with slightly different access protocol details. Separating the description of these two aspects helps WSDL represent common functionality between seemingly different end points.

Abstract Description

WSDL defines a service's abstract description in terms of messages exchanged in a service interaction. There are three main components of this abstract interface: the vocabulary, the message, and the interaction. Agreement on a vocabulary is the foundation of any type of communication. WSDL uses external type systems to provide data-type definitions for the information exchange. Although WSDL can support any type system, most services use XSD. Figure 5 shows two data types defined in XSD (`string` and `int`), and two data types defined in external schema (`Flight`

InfoType and Ticket). WSDL can import such external XSD definitions using an “import” element specifying their location.

WSDL defines `message` elements as aggregations of parts, each of which is described by XSD types or elements from a predefined vocabulary. Messages provide an abstract, typed data definition sent to and from the services. The example in Figure 5 shows the three messages that might appear during a Web services interaction. The message, `GetFlightInfoInput`, has two parts: `airlineName`, which is an XSD string, and `flightNumber`, which is an XSD integer. The other two messages, `GetFlightInfoOutput` and `CheckInInput` have only one part each. The `operation` and `portType` elements combine messages to define interactions. Each operation represents a message exchange pattern that the Web service supports, giving users access to a certain basic piece of service functionality. An operation is simply a combination of messages labeled as `input`, `output`, or `fault` to indicate what part a particular message plays in the interaction.

A `portType` is a collection of operations that are collectively supported by an end point. In our example, `AirportServicePortType` describes two operations: a single request-response operation, `GetFlightInfo`, which expects the `GetFlightInfoInput` message as input and returns a `GetFlightInfoOutput` message as the response; and a one-way operation, `CheckIn`, which just takes the `CheckInInput` message as input.

Concrete Binding Information

So far, all of the elements that we have discussed describe the service’s application-level functionality. To complete the description of client-service interaction, we need three more pieces of information:

- *what* communication protocol to use (such as SOAP over HTTP),
- *how* to accomplish individual service interactions over this protocol, and
- *where* to terminate communication (the network address).

WSDL’s binding element provides the “what” and “how” parts of this information, including the communication protocol and data format specification for a complete `portType`. In short, the binding element tells how a given interaction occurs over the specified protocol. Figure 6 (next page) shows a fragment from our example. The binding describes how to use SOAP to access the `travelservice` ser-

```
HTTP/1.1 200 OK
Content-Type: text/xml; charset="utf-8"
Content-Length: nnnn

<SOAP:Envelope xmlns:SOAP=
  "http://schemas.xmlsoap.org/soap/envelope/">
  <SOAP:Body>
    <m:GetFlightInfoResponse
      xmlns:m="http://www.acme-travel.com/flightinfo"
      SOAP:encodingStyle=
        "http://schemas.xmlsoap.org/soap/encoding/"
      xmlns:xsd="http://www.w3.org/2001/XMLSchema"
      xmlns:xsi=
        "http://www.w3.org/2001/XMLSchema-instance">
      <flightInfo>
        <gate xsi:type="xsd:int">10</gate>
        <status xsi:type="xsd:string">ON TIME</status>
      </flightInfo>
    </m:GetFlightInfoResponse>
  </SOAP:Body>
</SOAP:Envelope>
```

Figure 4. SOAP RPC response. The travel service responds to Joe’s request with a structured value containing subvalues for gate number and flight status.

```
<message name="GetFlightInfoInput">
  <part name="airlineName" type="xsd:string"/>
  <part name="flightNumber" type="xsd:int"/>
</message>

<message name="GetFlightInfoOutput">
  <part name="flightInfo" type="fixsd:FlightInfoType"/>
</message>

<message name="CheckInInput">
  <part name="body" element="eticketxsd:Ticket"/>
</message>

<portType name="AirportServicePortType">
  <operation name="GetFlightInfo">
    <input message="tns:GetFlightInfoInput"/>
    <output message="tns:GetFlightInfoOutput"/>
  </operation>
  <operation name="CheckIn">
    <input message="tns:CheckInInput"/>
  </operation>
</portType>
```

Figure 5. WSDL abstract description. This fragment shows the string and int data types, which are defined in XSD, and two other data types defined in external schema: `FlightInfoType` and `Ticket`, which we assume were imported earlier in the WSDL file.


```

<binding name="AirportServiceSoapBinding"
  type="tns:AirportServicePortType">
  <soap:binding transport=
    "http://schemas.xmlsoap.org/soap/http"/>
  <operation name="GetFlightInfo">
    <soap:operation style="rpc"
      soapAction="http://acme-travel/flightinfo"/>
    <input>
      <soap:body use="encoded"
        namespace="http://acme-travel.com/flightinfo"
        encodingStyle=
          "http://schemas.xmlsoap.org/soap/encoding/" />
    </input>
    <output>
      <soap:body use="encoded"
        namespace="http://acme-travel.com/flightinfo"
        encodingStyle=
          "http://schemas.xmlsoap.org/soap/encoding/" />
    </output>
  </operation>
  <operation name="CheckIn">
    <soap:operation style="document"
      soapAction="http://acme-travel.com/checkin"/>
    <input>
      <soap:body use="literal"/>
    </input>
  </operation>
</binding>

<service name="travelservice">
  <port name="travelservicePort"
    binding="tns:AirportServiceSoapBinding">
    <soap:address location=
      "http://acmetravel.com/travelservice"/>
  </port>
</service>

```

Figure 6. WSDL's concrete binding information. As this fragment shows, *GetFlightInfo* is a SOAP RPC interaction and *CheckIn* is a pure messaging interaction that uses XSD to describe the transmitted XML.

vice. In particular, the WSDL document shows that

- *GetFlightInfo* will be a SOAP-RPC-style interaction, in which all message exchanges use standard SOAP encoding, and
- *CheckIn* is a pure messaging interaction ("document-oriented," in WSDL terms) in which the SOAP message's body contains the encoded message with no additional type encoding; that is, it uses XSD to literally describe the transmitted XML.

All that remains now is to define "where" to access

this combination of abstract interface and protocol and data marshalling details (the binding). A port element describes a single end point as a combination of a binding and a network address. Consequently, a service element groups a set of related ports. In our travel service example, a single port describes an end point that processes SOAP requests for the *travelservice* service.

Using WSDL

For users and developers, WSDL provides a formalized description of client-service interaction. During development, developers use WSDL as the input to a proxy generator that produces client code according to the service requirements. WSDL can also be used as input to a dynamic invocation proxy, which can then generate the correct service requests at runtime. The result in both cases is to relieve the user and developer of the need to remember or understand all the details of service access. For example, travel service users need only obtain the WSDL description and use it as input to the tooling and runtime infrastructure to exchange the correct SOAP message types with the service.

Discovery: UDDI

The Universal Description, Discovery, and Integration specifications offer users a unified and systematic way to find service providers through a centralized registry of services that is roughly equivalent to an automated online "phone directory" of Web services. The browser-accessible UDDI Business Registry (UBR) became available shortly after the specification went public. Several individual companies and industry groups are also starting to use "private" UDDI directories to integrate and streamline access to their internal services.

UDDI provides two basic specifications that define a service registry's structure and operation:

- a definition of the information to provide about each service, and how to encode it; and
- a query and update API for the registry that describes how this information can be accessed and updated.

Registry access is accomplished using a standard SOAP API for both querying and updating. Here we focus on the first aspect, which provides a good idea of how the registry operates.

Organizing Structure

UDDI encodes three types of information about Web services:

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.