

Conceptual Overview of the Implementation of Web Services Specifications in FINS

Community Grids Laboratory
Indiana University, USA

In WS specifications the logic contained within the specifications and the protocol is achieved through the exchange of SOAP messages. The logic pertaining to the specification can be encapsulated in the header and/or the body of the SOAP messages. Individual SOAP messages are self-descriptive.

1. Function performed by Software

The FINS implementation provides Grid and Web Service applications with the ability to issue notifications to each other based on the implementation of the WS-Eventing specification. Nodes can register their interests with the services in a variety of subscription formats.

2. Design of the WsProcessor

In our approach to implementing the Web Service specifications we considered SOAP messages as the focal point. All implementations of the WS specifications (in FIRMS and FINS) extend the WsProcessor class that is part of the `cgl.narada.wsinfra` package. This WsProcessor contains just one method viz. `processExchange(SOAPContext, direction)` where SOAPContext simply encapsulates the SOAPMessage. Processors (either WSRM or WSE for example) process these SOAP messages based on the processing outlined in the respective specifications. Furthermore, in instances where these specifications leverage other WS specifications, such as WS-Addressing, the rules outlined in those specifications are also enforced. For example the construction of SOAP Envelope responses is based on the WSA message information headers (MIH) and also on the rules governing the mapping of WSA Endpoint References (EPR). Note that these processors provide complete implementations of the specification, i.e. they generate the appropriate responses, actions or faults in response to SOAP messages that have been received.

Included below is the definition of the `processExchange()` method. Using the SOAPContext it is possible for an entity to retrieve the `javax.xml.SOAPMessage` or the equivalent `EnvelopeDocument` (from XMLBeans). The direction specifies whether the message was received over the network or from the application. The logic related to the processing of messages is different depending on whether the message was received from the application or network. Exceptions thrown by this method are all checked exceptions and can be trapped using appropriate try-catch blocks. Examples of when these exceptions are thrown are described in the section outlining how the WsProcessors can be deployed within handlers.

```
public void processExchange(SOAPContext soapContext,
                           int direction) throws UnknownExchangeException,
                                                IncorrectExchangeException,
                                                MessageFlowException,
                                                ProcessingException
```




Figure 1: The WsProcessor and WsMessage flow components

Please note that in most WS specifications there are multiple roles. For example in WSRM the two distinct roles are that of a source and a sink. Similarly, in WS-Eventing the roles outlined by this specifications include source (notification producer), subscription manager and sink (notification consumer). A processor decides on processing a SOAP based on three parameters

- The contents of the WSA action attribute contained within the SOAP Header.
- The presence of specific schema elements in either the Body or Header of the SOAP Message.
- If the message has been received from the application or if it was received over the network.

If the WS processor does not know how to process a certain message, it throws an UnknownMessageException an example of this scenario is a WS-Eventing source node receiving a CreateSequence response from over the network. An IncorrectExchangeException is thrown if the WsProcessor instance should not have received a specific exchange. For example if a WSRM sink node receives a wsrn:Acknowledgement it would throw that particular exception. Please note that these exceptions are NOT fatal errors, they simply indicate what the WsProcessor instance (whether it is a WsrnSourceNode, WsrnSinkNode, WseSourceProcessor, WseSinkProcessor or WseSubManager). Furthermore, in most cases these WsProcessors will be cascaded together. If a user wishes to use both the Wsrn Source and Sink capabilities the corresponding WsProcessor instances need to be instantiated.

MessageFlowException and ProcessingExceptions are errors caused due problems with networking and processing a message respectively. Typically, when these exceptions occur unlike the previous exceptions processing related to the message within the handler/filter chain needs to be terminated immediately. ProcessingExceptions occur due to processing errors related to inability to locate protocol elements in message, incorrect schemas and no values being supplied for some elements.

2.1 Role of the WsMessageFlow.

With regards to deployment, the other class of interest is the WsMessageFlow. This interface contains two methods enrouteToNetwork(SOAPMessage, address) and enrouteToApplication(SOAPMessage). The WsProcessor uses these methods to route SOAP messages (requests, responses or faults) en route to applications or a network endpoint respectively. Since the WsProcessor delegates the actual transmission of messages to implementations of the WsMessageFlow, it can be deployed in a wide variety of settings, for example as a handler, proxy and within application programs. Note that, typically the job of transmitting messages is within the purview of the hosting environment viz. Axis, JWSDP etc.

2.2 Rationale for the choice of XMLBeans

While implementing these specifications we were faced with an important decision regarding the choice of tool to use in processing the XML schema that aforementioned specifications conform to. In simple terms we were looking for a system that allowed us to process XML for within the Java domain. There were three main choices. First, we could use the AXIS wsdl2java compiler. It is well-known that Axis' support for complex Schema Types is poor. In fact, the XML generated by the generated Java classes can sometimes not be compliant with the original schema specification. Finally, support for the XML schema in Axis is both buggy and incomplete. The problems that we outlined with Axis also exist in Sun's JWSDP. We need true support for the XML Schema since at no point can we make the assumption that we will be dealing with Axis/JWSDP endpoints, which would conform to a subset of the XML schema.

The second approach was to use the JAXB specification (a specification from Sun to deal with XML and Java data-bindings). JAXB though better than what is generated using Axis' wsdl2java still does not provide complete support for the XML Schema. We looked at both the JAXB reference implementation from Sun and JaxMe from Apache (which is an open source implementation of JAXB). We note that this while significantly better than Axis' wsdl2java still does not provide complete support for the XML schema.

The final approach involves utilizing tools which focus on complete schema support. Here there were two candidates XMLBeans and Castor which provide good support for XML Schemas. We settled on XMLBeans because of two reasons. First, it is an open source effort. Though originally developed by BEA it was contributed by BEA to the Apache Software Foundation. Second, in our opinion it provides the best and most complete support for the XML schema of all the tools currently available. It allows us to validate instance documents and also facilitates simple but sophisticated navigation through XML documents. The XML generated by the corresponding Java classes is true XML which conforms to (and can be validated against) the original schema.

3. Specifications and schemas used

As most folks who have tracked WS specifications would agree, there are sometimes more than 1 version of the specification and the concomitant schemas. During the implementation of the WSRM and WS-Eventing specifications we made a decision to use the latest schemas. In some cases (such as SOAP) we were forced to use the older schema available at <http://schemas.xmlsoap.org/soap/envelope/>, since the javax.xml.SOAPMessage is based on this one. We do not however foresee problems in migrating to a newer version of the SOAP schema based on SOAP 1.2 should a need arise. We implemented the latest version of the WS-Eventing specification so the schema corresponds to this version. We enumerate below the schemas corresponding to the various WS-Specifications and XML utilities that we used in our implementations.

FINS provides an implementation of the WS-Eventing specification that was released in August 2004. This specification developed jointly by Microsoft, BEA, IBM and Sun can be found at <http://www-128.ibm.com/developerworks/webservices/library/specification/ws-eventing/>
 WS-Eventing: <http://schemas.xmlsoap.org/ws/2004/08/eventing>
 SOAP: <http://schemas.xmlsoap.org/soap/envelope/>
 WS-Addressing: <http://schemas.xmlsoap.org/ws/2004/08/addressing>

4. Deployment and other Issues

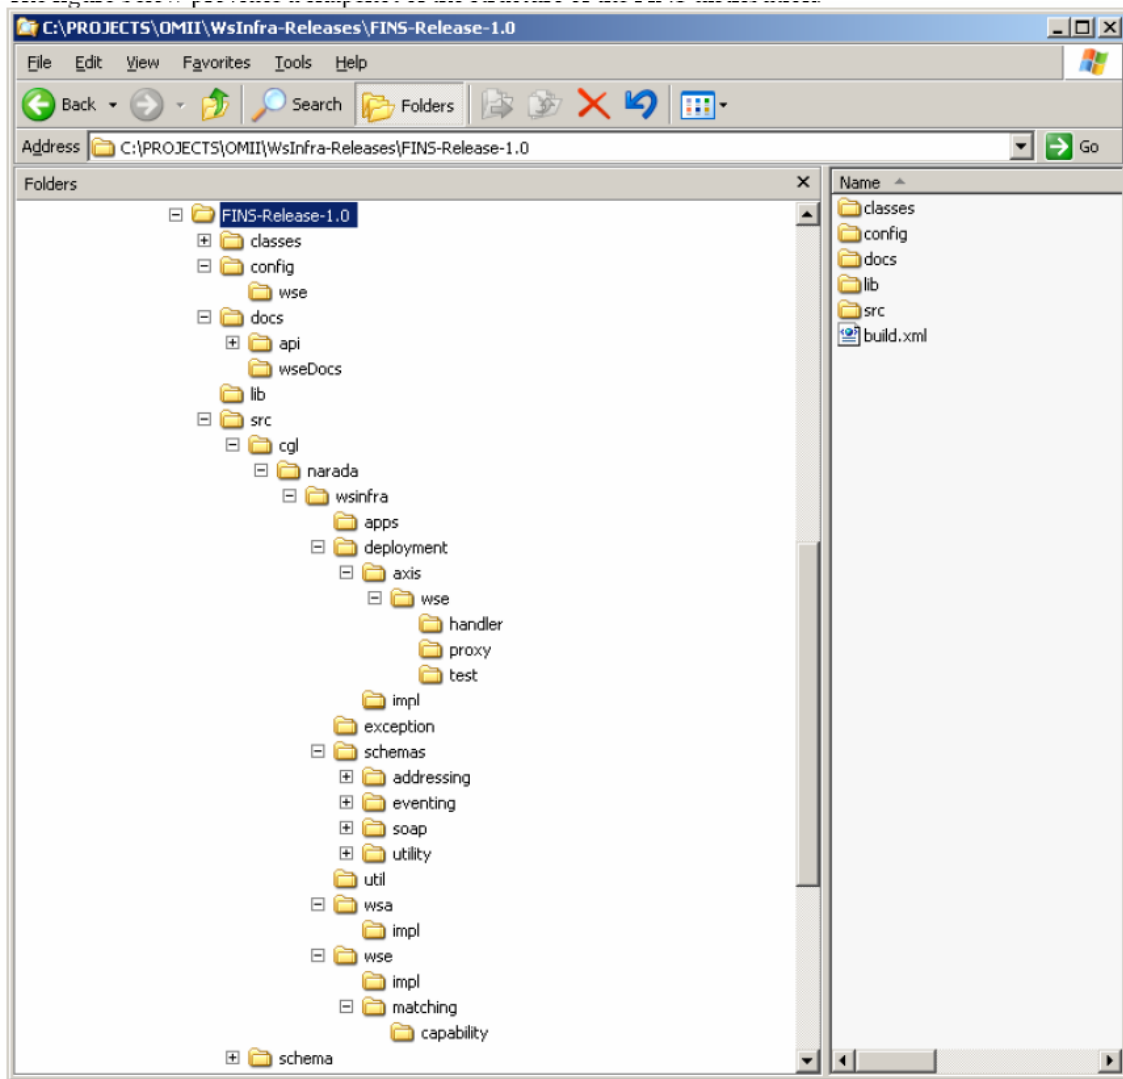
This release includes two different deployments of the implementation of the WS specifications. The first deployment environment is the OMII Container (we have successfully deployed the software within the OMII 1.2 container). The OMII container leverages the Axis Web Services container. Please see document in the docs directory for a descriptions of problems related to Axis and the strategies that we have incorporated to circumvent these problems. To facilitate testing of the WS-Eventing software without having to cope with the limitations of the deployment containers (as discussed in a meeting with the OMII in January 2005) we built a prototype deployment container. This prototype deployment container has been included in this release and is based on one-way SOAP messages. The prototype deployment container is based on Filter/Filter-Pipeline architecture. Some of its features include

1. Dynamic configuration of the Filter Pipeline. Filters can be added, removed, updated or reorganized within the Filter-Pipeline.
2. Individual filters can configure themselves to be part of Input, Output or Input/Output processing.
3. Filters can inject messages either towards the application or towards the network. The Filter-Pipeline ensures that appropriate filters are traversed depending on the direction of traversal.
4. A filter can terminate processing related to a message within the Filter-Pipeline.
5. A filter can add, remove, modify or replace the contents of the SOAP Message or the equivalent EnvelopeDocument.
6. It is possible to configure different networking substrates with the Filter-Pipeline thus delegating the networking capabilities to it. Destinations are computed based on WS-Addressings [wsa:To] element.
7. The filter model and its capabilities such as injecting messages, terminating processing and the ability to be cascaded provide a true representation of how WS specifications are intended to provide incremental addition of capabilities at an end-point.

Please note that the primary motivation for writing this prototype deployment container was to have a test environment that allows us to test every exchange as outlined in the implemented Web Service specifications and see if the intended action corresponding to each exchange is taken and faults issued if necessary.

5. Package Structure

The figure below provides a snapshot of the structure of the FINS distribution.



6. Future Releases

Future releases will include open-source implementations of the WS-Notification suite of specifications. Please track the appropriate web sites for information regarding updates.