

# Automating Metadata Web Service Deployment for Problem Solving Environments

Ozgur Balsoy<sup>a,b</sup>, Ying Jin<sup>b</sup>, Galip Aydin<sup>b</sup>, Marlon Pierce<sup>b</sup>, and  
Geoffrey Fox<sup>b</sup>

<sup>a</sup>*Department of Computer Science, Florida State University,  
Tallahassee, FL 32306*

<sup>b</sup>*Community Grids Lab, Indiana University,  
501 N.Morton Stree, Suite 224, Bloomington, IN 47404-3730*

---

## Abstract

XML-based metadata information services are a crucial core service needed by Problem Solving Environments built over emerging service-based, globally-scaled distributed systems, as envisioned by the Open Grid Services Architecture and the Semantic Web. Developing user interfaces and services bindings for manipulating instances of particular schemas is thus extremely important and needs to be made as simple as possible. In this paper we describe procedures for automating the creation of Web Service environments that can be used to simplify the creation and deployment of schema-based metadata services.

*Key words:* Metadata authoring, Problem Solving Environments, XML, Web Services

---

## 1 Introduction

One of the most important underpinnings of Problem Solving Environments (PSEs) is the management of diverse types of information. Grid computing and distributed Web services are metadata rich: all entities should be described by XML metadata. Metadata must be used to describe the various parts that must be managed by the PSE, such as available applications, service interfaces and their binding points, batch queuing systems, hardware profiles, and so on. Many projects have attempted to define metadata pieces needed to build service grids, and such pieces will be needed as grid service data in the Open Grid Service Architecture [1,2]. In addition to the service metadata, there

is also need to provide organizational metadata descriptors for the services themselves, and we have suggested data models for describing applications [3,4] and their associated service bindings. We thus anticipate that metadata richness will only increase.

We may place these issues into larger contexts. *Autonomic computing* (see [5] for a description and [6,7] for sample academic projects) seeks to define the interaction of components in next generation computing at all scales. In the autonomic computing vision, components mean everything from the physical component parts of individual computers to the service components of globally-scaled service environments. The ideal autonomic system is self-aware, self-monitoring, adaptive, and self-healing, among other characteristics.

Developing standards for Web services for computational grids is a major (and somewhat independent) component of the autonomic computing vision. Web services are somewhat loosely defined, but may be characterized in general as using existing Web technologies and standards (such as HTTP and XML) to build the distributed computing environments. Important Web service standards include the Web Service Description Language (WSDL) [8] and the Simple Object Access Protocol (SOAP) [9]. The Open Grid Service Architecture, developed under the auspices of the Global Grid Forum and primarily led by IBM and the Globus team, seeks to extend the capabilities of Web services in order to support more sophisticated services. The key to this development is the extension of WSDL portTypes to include metadata and service state information.

In a major alternative effort, other groups are using various XML standards of the World Wide Web Consortium define what it terms the Semantic Web [10]. The Semantic Web vision borrows ideas from federated information systems and artificial intelligent, modified and simplified to apply to the much more ambitious problem of intelligent information retrieval and service execution on a global scale. The cornerstones of the Semantic Web are metadata descriptions (RDF) [11] and ontologies that provide semantic meaning.

XML-based information services are the key to both the OGSA and the Semantic Grid. We may go a step further and claim that there is a general need for many informal, lightweight information management systems for particular schemas. We have developed such a system, described in Refs [12] and [13], which provides several per-schema implementations: newsgroups, citation and reference managers, training and conference registration systems. All such systems possess the same general capabilities: browsing, search, "push" and "pull" style event delivery options, and access controls on information and publishing. More sophisticated systems that support schema-spanning searches (with the associated meta-ontology research problems [14,15] may be built on top of this practical base.

Given the demonstrable importance and proliferation of XML metadata, it becomes necessary to automate the creation of publication and editing tools for particular schemas. We describe the details of two such systems, which we refer to as "Schema Wizards". The first is designed to simplify the creation of XML wizards for specific schemas. These wizards are in turn used for creating and publishing instances in the XML messaging system described in [13]. The Schema Wizard is thus a wizard for wizards. The second wizard system, a follow-on to the first, is designed to automate metadata-based Web services (such as service data associated with Grid Web services). In both cases, a particular schema is used as the underlying data model, and the wizard generates the user interface (view) and action bindings (control) necessary to automatically create and deploy the new metadata-based service.

## 2 Schema Wizard: Automating Metadata Creation for Publication

Applications that guide users through a complicated task are usually called wizards. XML wizards are applications that receive user inputs to generate XML documents. In our research on XML information frameworks (see particularly Ref [13], XML is used as the medium for storing information, but the systems are user driven. Wizards help users generate schema-based XML content and publish into an XML messaging broker. Such documents are then typically published to various sources, including both persistent repositories and listeners demanding immediate notification.

In our current implementation, our information management system provides general purpose tools that can be used to create specific applications that are tied to single, specific schema. Thus to add a new application, one of the necessary tasks is to create a new publishing wizard to the particular schema. Creating such publishing wizards requires the following steps:

- (1) Writing an XML Schema which has all the necessary elements regarding with the current project.
- (2) Producing the Object Models of the XML Schema elements and attributes. This is done by using Castor's [16] Source Generator, which creates Java beans to correspond to the elements of the schema.
- (3) Obtaining user data. This step requires us to prepare HTML forms which have form elements to represent the schema elements. We use JavaServer Pages (JSP) technology, which allows us to interact with HTML forms and data objects at the same time. In that way we immediately assign the user entries to the data objects.
- (4) Generating the XML document. Using Castor's marshalling framework we create an XML document which contains the user inputs.
- (5) Publishing this XML document. This may use publish/subscribe systems




Fig. 1. The Schema Wizard is used to generate interfaces for particular XML schemas, which in turn generate XML instances for messaging.

such as the JavaMessaging System (JMS).

Given the repetitive, tedious nature of this process, we have developed a general purpose tool to that automates the creation of this process. We refer to this tool as the SchemaWizard.

### 2.1 Schema Wizard Architecture

The Schema Wizard (SW) maps XML Schema elements to HTML form elements through its schema parser, and creates the framework and logic for an XML form wizard, as depicted in Figure 1. The Schema Wizard is used to generate interfaces for particular XML schemas, which in turn generate XML instances for messaging.

Users can use newly generated wizards to create, edit, and publish XML instances, which follow a given schema, to any destinations such as publish/subscribe messaging systems or through SMTP. XML form wizards are Web applications that can also serve as validating XML editors and be customized through schema annotations.

The SchemaWizard has been developed as a Web application compliant to JSP/Servlet specifications that generates and deploys new Web applications to help instantiate XML documents. A SW user defines her data model using a XML schema and annotates the schema to customize the HTML form view. [Fig. 2] After provided with a schema, the Wizard first creates the necessary directory structure and required libraries for a XML form wizard by unpacking a Web application template package. Second, the Wizard invokes Castor SourceGenerator [16] to create Java object definitions (a.k.a. javabeans) corresponding to individual XML schema elements, and compiles them into the previously created Web application directories. As the third step, the Wiz-




Fig. 2. The XML Schema Wizard simultaneously generates user interfaces, in-memory data bindings, and actions links between interface elements and data objects for an XML schema.

ard uses Castor’s Schema Object Model (SOM) API to parse and process the schema in the memory. The SOM API is used for its convenient and specialized schema type access methods, which we prefer for this purpose over other XML processing APIs, i.e. DOM or JDOM.

The SchemaParser module, the heart of the Wizard application, traverses the schema using the SOM API and collects XML type and structure information. This information later is used to decide what type of template is necessary to generate final JSP pieces or *nuggets*. Velocity [17] templates are used for their simplicity to define JSP nuggets, each of which maps either a simple schema type or a complex type to one or more HTML elements rendered by JSP engine. The JSP page created for the root element is included by a higher level index page where the entire form view is finally constructed. The rules for mapping from XML schema types to HTML elements are listed in the following section.

## 2.2 Schema Element to HTML Mapping Rules

To generate XML instances based on XML schema elements, each schema element is mapped to one or more corresponding HTML form elements. Mapping follows these guidelines:

- (1) One HTML form page is generated for each possible root (or global) element, which is based on an anonymous or named global complex type.
- (2) XML schemas define elements for any level within XML document hierarchy. Some elements, for example, may not need any parent element to exist and can root a XML document. Such elements are called global elements and defined at the top level in a XML schema. Elements can also include sub elements and attributes. These elements are based on

- complex type definitions. SW generates forms for only global elements that are based on either named or anonymous complex types.
- (3) Each form has one form level submission, which submits all the content inside form elements and alters current values on the server side at once. Each form may also have more than one element level submission, which affects only related elements or attributes.
  - (4) Simple, non-enumerating, non-repeating elements or attributes are mapped to a single input text field.
  - (5) Simple typed elements or attributes must be based on predefined XML schema types such as numbers, character strings, and date and duration types. Any simple type can be restricted into a narrower type range. For an example, a new type can be generated from the integer type by restricting the values to the integer numbers between 0 and 100. Another example could be a character string type that must follow a pattern such as social security numbers with 9 digits and two dashes. Such schema types are mapped by SW to single input text fields and their values are only updated if a form level submission occurs. Data validation is performed on the server side while forms are used for some input size restrictions.
  - (6) An enumerated simple type is mapped to a selection box with the enumerated values filled in. XML schemas allow some of the simple element or attribute values to be enumerated. Enumerated values are provided with type definitions inside schemas. Drop-down selection boxes with enumerated values filled in are used to map such types. If an enumerated element or attribute is optional, users are provided with an option to remove it. These types are updated by only form level submissions.
  - (7) An unbounded simple type is mapped to an input field, a selection box, and two buttons to add and remove new values in addition to form level submissions. Unbounded simple types are elements that can repeat with the same or different values. The content of the input field is submitted and a new element is created when the submission is performed through the addition button. The new element is also added to the selection box. Users can remove any element by choosing its value from the selection box and submitting the form through the remove button.
  - (8) A defaulted simple type is mapped to an input field with a default value.
  - (9) A fixed value simple type is mapped to no form element, but the fixed value is written in plain text.
  - (10) A simple Boolean-typed element or attribute is mapped to a pair of radio buttons with labels, "Yes" and "No."
  - (11) An unbounded complex value is mapped to, along with its sub elements, a selection box where the element's index can be chosen, and a pair of add and remove buttons. If any complex element index is chosen from the list, it can be removed by submitting through the remove button.

### 2.3 View Construction

The following is a sample Velocity template for simple enumerated types followed by the JSP page generated from this template. All the Velocity variables starting with \$ and surrounded by braces are replaced with the schema types information gathered by the SchemaParser. Elements and attributes are treated the same way in terms of the interface however, Castor SourceGenerator generates different source code for them, and the template reflects this through a Velocity variable, \$attribute. Also, enumerated values are embedded into the JSP nugget through an array of values, \$values, obtained through the SOM API and forwarded to the Velocity template engine by SchemaParser.

```
<% {
  // obtain element's name and its submitted value
  String elementName = "$name";
  String value = request.getParameter(elementName);

  // obtain if this element has a parent. If root, no parent.
  Stack parents = (Stack)request.getAttribute("_parents_stack");
  String parent = (parents==null —— parents.empty())
    ? null
    : (String) parents.peek();

  // check if form is submitted by either the main submit button,
  // or other buttons on the form, such as those
  // for unbounded elements
  if(request.getAttribute("submit_parent") != null
    && value != null)
  {
    // if [remove] item is selected, remove this element.
    // otherwise set its value
    if(value.equals(REMOVE))
      ${parentname}.set${javaname}(null);
    else
      try {
        ${parentname}.set${javaname}
          (${javaname}Type.valueOf(value));
      } catch(Exception ex) {}
  }
  else if(parent != null)
  {
    // if the form is not submitted by the main submit button,
    // then bring back the in-memory value
    Object valueObj = ${parentname}.get${javaname}();
```

```

    value = (valueObj==null) ? null : valueObj.toString();
}
if(value == null)
    value = "";

// enumerated values in an array
String[] values = {
#set ($i = 0)
#foreach ($enum in $values)
"$enum" #set ($i = $i + 1) #if( $i < $values.size() ) , #end
#end
};

// send the HTML for enumeration
%>
<tr>
<td>
## if numbers is requested, then print out the leading numbers,
## i.e. 1. , 2.2, 3.4.2
#if ($numbers)
<%= ((Leveler)request.getAttribute("_leveler")).increase()%>
#end
$label</td>
<td>
<select name="$name">
    <option value="<%=REMOVE%>"><%=REMOVE%></option><%=
// print out option. select the one that matches in-memory value
for(int i=0; i< values.length; i++)
{
    String enm = values[i]; %>
    <option value="<%= enm %>"<%=
        (value != null && value.equals(enm))
        ? " SELECTED"
        : "" %>><%= enm %>
    </option><%=
} %>
</select>
</td>
</tr><%=
} %>

```

An example of the JSP generated from this template may be found in [18]. The parent element of this type is a complex type called `rootElement`. The JSP

starts with retrieving the user input data, and checks whether this element has a parent or not. Based on the result of this check, the JSP nugget decides whether the memory object needs to be updated or the form is to be filled with the old data in memory. Finally, a selection box is printed out in HTML with enumerated values stored in an array of strings, values[].

All variables defined in generated JSP nuggets are local. Since parent elements include sub elements' JSPs, initializations are performed within parents' scope. The following is an example generated JSP for a parent element. The last part of the JSP decides whether this element is the root element and the submission buttons need to appear.

In XML schema, all the elements defined as global elements can be referenced at any level within the schema hierarchy. To support referenced global elements, we have designed the JSPs for parent elements, such as the example above, so that they can be included by other complex typed elements. Here, the complex type 'rootElement', can repeat in any part of the schema as well as in its mapped HTML form.

The entry pages for HTML forms are index pages, which must not be included by other JSPs. SW does not provide such inclusions. Developers who use generated JSPs should also notice this requirement.

### **3 WSDL Wizards for Metadata Web Services**

In the proceeding section we described an XML Schema Wizard that can be used to automate the process for creating specific wizards for generating XML messages. This system is appropriate for publish/subscribe style interactions. Here we describe our design for a variation of this system that is used to automatically generate remote procedure call (RPC)-style services. Instead of supporting all schema types, we need only to support one: WSDL. User interfaces generated from WSDL instances will be used to create server-side SOAP calls to back-end XML repositories.

As we argued in the introduction, an important subclass of Web services must be geared toward the creation and management of XML metadata. We refer to this particular subclass of Web services as Metadata Web services (MWS). Creating and deploying a new MWS has the following steps:

- (1) Create Java class bindings for the schema.
- (2) Create a Faade class to serve as Web service implementation and simplify manipulation of the data classes.
- (3) Deploy the service.

- (4) Generate client stubs from the service's WSDL.
- (5) Create the user interface to the service using the client stubs.

The Faade class for the server (step 2) may be used to simply manipulate the schema data objects, but for complicated schemas it also serves to simplify (at the expense of lost functionality) the task of step 5. These steps are somewhat tedious to perform manually. The tedious nature of creating these services has some unpleasant side effects: it inhibits the modification phase of trial schemas and (more globally) discourages the proliferation of lightweight, schema-based services.

The Schema Wizard itself, as described above, is designed for creating XML messages. In particular, it is not directly suitable for creating services and user interfaces from the WSDL schema. The WSDL document consists of two parts: service interface definitions and implementation bindings. The interface part contains the abstract definition of a service, which allows one to retrieve information of operations, and the types of the messages the operations receive and return. The implementation section contains concrete information on port and binding protocol. We thus do not want to create the user interface for a WSDL service description from the entire document but instead only the abstract portion (the portTypes) of the document.

The WSDL Wizard is designed to automate all parts needed to create a Metadata Web Service. This essentially has the following parts: a) selection of desired WSDL instance, b) generation of user interface and server-side bindings, and c) dynamic invocation of WSDL interface methods to manipulate the remote XML instance.

We are developing a Java package that allows SOAP procedures to be dynamically constructed for a given WSDL instance without needing to create client stub programs. This is described elsewhere [19], and in this paper we focus its combination with the Schema Wizard.

The major components need to accomplish this are illustrated in Figure 3. The user is assumed to have selected a particular service by inspecting a WSIL [20] document through a WSIL proxy client. The WSDL description of this interface is downloaded and the service generation procedure begins.

The important additional components are the WSDL Parser, GhostWSDL Parser, and User Data Parser. These parsers each are designed to provide access for manipulating in-memory representations of XML. A Controller component is responsible for managing all the components. The WSDL Parser is given the WSDL document that describes the desired service and creates an in-memory representation of the service interface. This document is then inspected by the Controller in order to dynamically generate a SOAP call.



Figure 3 shows how these components are related to each other. The arrow lines indicate the communication between the two components. The diagram uses the thicker lines for Internet request/response messages with the indicated protocols, and the thinner lines represent internal communications. The following is detailed explanation for each numbered step.

- (1) The user requests the service list through a browser interface.
- (2) The JSP interface sends the user's request to Controller
- (3) Controller communicates with the web services inspection proxy client (a WSIL client).
- (4) Proxy client makes the SOAP call to the web services inspection proxy server and requests the WSIL service list.
- (5) Proxy server makes the HTTP connection to a remote/local server for getting WSIL document
- (6) The remote/local server sends the WSIL document back to the proxy server
- (7) Proxy server sends back the service list to the client
- (8) The proxy client responds to Controller
- (9) Controller returns the service list to the user
- (10) The user picks a desired Web service from the list WSIL list.
- (11) Steps (2) - (8) are repeated in order to retrieve the WSDL document, this time going through (5) and (6).
- (12) The WSDL document is sent to the parsers. In 12a, the WSDLPARSER creates in memory representation of WSDL. In 12b, the GhostWSDL is extracted from WSDL instance.
- (13) The GhostWSDL is sent to the Schema Wizard, which generates the user interface.
- (14) The user enters the required information of service into the system.
- (15) The wizard generates a GhostWSDL instance with the user's data, is created. If the GhostWSDL interface contains complex types, then JavaBean classes are also generated.
- (16) A Java Object array with the user data is passed to Controller.
- (17) The controller obtains an in-memory representation of the WSDL instance.
- (18) The Controller inspects the WSDL document, constructs, and launches the SOAP call to invoke the service.

#### **4 Summary and Future Work**

As part of our attempts to address a larger metadata management problem, this paper has presented two mechanisms for automatically deploying web-based applications for creating and manipulating metadata based on XML schemas. The Schema Wizard is a general purpose system for creating user

interfaces and action bindings for a particular XML schema. As we have illustrated, the SW is used to generate browser interfaces automatically from a particular schema data model. It also binds the schema to in-memory data objects that can be used to manipulate schema instances. In addition to generating the view of the data model, we also generate the action bindings (in this case, HTML Form actions) that connect the user-supplied values to the data model. This is a useful component in message-based systems such as we describe in [13]. Future work on the Schema Wizard will be to introduce more customization of the presentation.

In contrast to one-to-many, message-style metadata services, the one-to-one, RPC-style metadata services are more meaningfully described with WSDL interfaces than message formats. To automatically generate the user interface to these services, we need to modify the Schema Wizard operation to display not the entire WSDL schema but the only the portion appropriate for display. In addition, the bindings must now change so that we can tie user interface actions to backend (rather than server-side) data instances, which we do with SOAP invocations dynamically generated from the WSDL instances. We have completed development of the preliminary libraries for dynamic WSDL invocations (supporting XML primitive types), and the WSDL wizard implementation based on these libraries is currently in progress. The next step will be to add support for WSDL complex types, as we describe in the paper.

## References

- [1] I. Foster, C. Kesselman, J. M. Nick, S. Tuecke, The physiology of the grid. an open grid services architecture for distributed systems integration, Open Grid Service Infrastructure WG, Global Grid Forum (June 2002).
- [2] S. Tuecke, K. Czajkowski, I. Foster, J. Frey, S. Graham, C. Kesselman, Grid service specification, Open Grid Service Infrastructure WG, Global Grid Forum (July 2002).
- [3] M. Pierce, G. C. Fox, C. Youn, S. Mock, K. Mueller, O. Balsoy, Interoperable web services for computational portals, in: Proceedings of High Performance Networking and Computing SC'02, 2002.
- [4] C. Youn, M. Pierce, G. Fox, Building problem solving environments with web services, in: Workshop on Complex Problem Solving Environments for Grid Computing, 2003.
- [5] Autonomic Computing: IBM's Perspective on the State of Information Technology (2001).  
URL  
[www.research.ibm.com/autonomic/manifesto/autonomic\\_computing.pdf](http://www.research.ibm.com/autonomic/manifesto/autonomic_computing.pdf)

- [6] S. Rhea, C. Wells, P. Eaton, D. Geels, B. Zhao, H. Weatherspoon, J. Kubiawicz, Maintenance-free global data storage, *IEEE Internet Computing* 5(5), 40-49. .
- [7] D. Abramson, R. Buyya, J. Giddy, A computational economy for grid computing and its implmentation in the nimrod-g resource broker, *Future Generation Computer Systems* 18.
- [8] E. Christensen, F. Curbera, G. Meredith, S. Weerawarana, Web service description language (wsdl) version 1.1.  
URL [www.w3c.org/TR/wsdl](http://www.w3c.org/TR/wsdl)
- [9] D. Box, D. Ehnebuske, G. Kakivaya, A. Layman, N. Mendelsohn, H. F. Nielsen, S. Thatte, D. Winer, Simple object access protocol (soap) version 1.1., W3C Note (May 2000).  
URL [www.w3.org/TR/SOAP](http://www.w3.org/TR/SOAP)
- [10] T. Berners-Lee, J. Hendler, O. Lassila, The semantic web, *Scientific American* (May 2001).
- [11] O. Lassila, R. R. Swick, Resource description framework (rdf) model and syntax specification, W3C Recommendation (1999).
- [12] O. Balsoy, et al., The Online Knowledge Center: Building a component based portal, in: *The Proceedings of Information and Knowledge Engineering (IKE)*, 2002.
- [13] G. Aydin, et al., An XML based system for dynamix message content creation, delivery, and control, in: *The Proceedings of IASTED International Conference on Information and Knowledge Sharing (IKS)*, November 2002.
- [14] J. Kahng, D. McLeod, Dynamic classification ontologies. computing the brain: A guide to neuroinformatics, Academic Press, 241-254 (2001).
- [15] S. Thacker, A. Sheth, S. Patel, Complex relationships for the semantic web, In *Creating the Semantic Web* (D. Fensel, J. Hendler, H. Liebermann, and W. Wahlster, eds). MIT Press (2001).
- [16] M. Birbeck, et al., *Professional XML*, 2nd Ed. (671-721), Wrox Press, Inc., 2001.
- [17] Apache jakarta-velocity project.  
URL [jakarta.apache.org/velocity/index.html](http://jakarta.apache.org/velocity/index.html)
- [18] O. Balsoy, Schemawizard project documentation (2003).  
URL [ptlportal.communitygrids.iu.edu/schemawizard/index.html](http://ptlportal.communitygrids.iu.edu/schemawizard/index.html)
- [19] Y. Jin, Web service dynamic invocation, Tech. rep., Community Grids Lab, Indiana University (2003).  
URL [grids.ucs.indiana.edu:8035/doc/Dynamics\\_Web\\_Services\\_Accessing.doc](http://grids.ucs.indiana.edu:8035/doc/Dynamics_Web_Services_Accessing.doc)
- [20] Web services inspection language (WSIL).  
URL [www-106.ibm.com/developerworks/webservices/library/ws-wsilspec.html](http://www-106.ibm.com/developerworks/webservices/library/ws-wsilspec.html)

## 5 Authors' Biographies



Ozgur Balsoy is a Ph.D. candidate at Florida State University. He has worked with Dr. Geoffrey Fox through his Master's and Ph.D. years working on and developing new technologies for the Internet and its effective usage in education. Currently, his research interests include metadata modelling and management for distributed and collaborative systems, metadata authoring tools, and component-based portal systems. He received his B.S. in computer engineering from Ege University, Turkey, and M.S. degree from Syracuse University.



Ying Jin has more than six years of software development experience, including network, database, and web design and development. Before entering IU as a graduate student, Ying worked in industry, focusing on printer driver implementation and Oracle database application development.

Right now, Ying is a Graduate Assistant in Pervasive Technology Labs at Indiana University. Her work involves processing research related to the Network Intrusion Detection System and Layer 2 Query Protocol.



Galip Aydin is a Ph.D. student in Computer Science Department and research assistant in Pervasive Technology Labs at Indiana University. He has built several XML wizards for various projects. His current focus is development of GML Schemas and Web Services for seismic events and fault data.

He has a B.S. degree from Electrical and Electronics Engineering at Firat University, Turkey. He had his M.S. degree in Computer Science from Syracuse University, NY in 2000.



Marlon Pierce is a Senior Postdoctoral Research Associate for the Community Grids Laboratory at Indiana University. His research interests include the development of Grid/Web services, component-based portal environments, and information representation and management to support scientific applications and data. He received a Ph.D. in Physics from Florida State University in 1998.



Geoffrey Fox, nationally renowned for his work in the development and application of parallel computers, most recently served as director of the Computational and Information Science Laboratory at Florida State University. He was also the director of the Northeast Parallel Architectures Center at Syracuse University from 1990-2000. His research has led to two commercial spinoffs-WebWisdom.com and Anabas Inc.

Fox's current projects include developing the Online Knowledge Center for the Department of Defense High Performance Computing and Modernization Program, which is creating a peer-to-peer system to allow users to more easily access and update information on the department's high performance computers.