

A NOAA Web Services Registry

FY 2004 Proposal to the NOAA HPCC Program

August 11, 2002

| [Title Page](#) | [Proposed Project](#) | [Budget Page](#) |

Principal Investigator: **Ann Keane**

Line Organization: ETL
Routing Code: R/ETL
Address:
NOAA/ETL
325 Broadway
Boulder, CO 80305

Phone: (303) 497-7424
Fax: (303) 497-6978
E-mail Address: Ann.Keane@noaa.gov

Eugene F. Burger	Eric A. Kihn	Mark W. Govett	Jeff Smits
NOAA/PMEL	NOAA/NGDC	NOAA/FSL	NOAA/OMI
7600 Sandpoint Way	325 Broadway E/GC2	325 Broadway	325 Broadway
Seattle, WA 98115	Boulder, CO 80305	Boulder, CO 80305	Boulder, CO 80305
Eugene.Burger@noaa.gov	Eric.A.Kihn@noaa.gov	Mark.W.Govett@noaa.gov	Jeff.Smits@noaa.gov

Proposal Theme: **Technologies for Collaboration, Visualization, or Analysis**

Ann Keane
Webmaster
NOAA/ETL

Rich Beeler
Chief, Computing and Network
Divison
NOAA/ETL

A NOAA Web Services Registry

Proposal for FY 2004 HPCC Funding

Prepared by: Ann Keane

Executive Summary:

Web services are beginning to provide an elegant solution to many of NOAA's data access challenges. We propose to establish a standards-based, web service directory or UDDI to serve as a public registry of web services and web service providers within NOAA. The directory will serve as a public registry for NOAA web services, allowing NOAA developers to register existing web services, search for registered web services and locate application programming interfaces (APIs) to access web service resources. The public will similarly be able to search for web services and access NOAA's web services through the published APIs. A central NOAA Web Service Registry will help developers to coordinate efforts, avoid redundant efforts and add momentum to the use of web service technology in NOAA. Several cross-Line Office web service development efforts will serve as the first beta-testers. We will advertise the availability of the NOAA Web Services registry through the NOAA CIO chain and at NOAA-wide meetings such as NOAA Tech and the NOAA Webshop and assist developers to use the UDDI.

Problem Statement:

Crontab and ftp are the duct tape of NOAA.

NOAA faces the challenge of disseminating its vast holdings of real-time and historical data to meet rapidly changing National needs. Access to data can be a costly, both in time and resources. Currently, data services require a great deal of human interaction to process, subset and distribute data. Even when data services are automated, they are not easily locatable or accessible by computer clients. Once data is located and routines are developed to automate access to the data, changes to the data format or distribution method can require rewriting the entire data access chain. This chain can also cause delays impacting real-time applications and operational reliability. Customized data access also increases the time needed to develop new applications and data products and slows the NOAA's response to the demand for new environmental products and services.

In order to meet this challenge, a number of NOAA groups have begun using Web Services technologies to distribute data. Web Services are a group of technologies developed and used commercially to automate access to data and computer resources and are beginning to be applied to environmental data services. While web services show promise for improving access to NOAA data and services, developers are working independently with little opportunity for interaction. One of the greatest strengths of web services is the ability to build modular services which can be used to create larger applications. Working in coordination, developers can use

existing web services to leverage new ones, increasing the efficiency and reliability of the new service and decreasing development costs. Coordinated web service development increases the efficiency of developers and prevents redundant efforts. Fostering the integration of web services development throughout NOAA is crucial to their success.

Although web services provide an elegant solution to many of NOAA's data access challenges, unlike web pages, webservices are not "linked" to one another, neither are they usually human readable. This makes publicizing web services a challenge. As the number of web service providers within NOAA increases, the need for some type of registry becomes more apparent. NOAA web services stand to be underutilized or duplicated if there is not a mechanism for registering, searching and accessing those resources. A central web service registry would add momentum and value to the use of web service technology in NOAA.

In summary, NOAA data and service providers are starting to develop web services as a solution to our data exchange requirements. As web services become more common in NOAA, there will be an increasing need for registry services including browsing and searching for existing web services. Without coordination, redundant efforts will fail to maximize the benefits of web services.

Relationship to HPCC objectives

This proposal falls under the "Technologies for Collaboration, Visualization, or Analysis" category. This proposal addresses the HPCC goal to "*explore the use of leading-edge E-government framework technologies such as Extensible Markup Language ([XML](#)) and standard schemas such as Electronic Business using eXtensible Markup Language ([ebXML](#)), Web Services Description Language ([WSDL](#)), Simple Object Access Protocol ([SOAP](#)), and Universal Description, Discovery, and Integration ([UDDI](#)) within NOAA are encouraged.*"

This proposal also allows "*for the inclusion of the broadest range of potential users*", as it seeks to foster NOAA-wide webservices resources and coordinate cross-line office efforts to integrate development of common resources.

Proposed Solution:

Synopsis

We propose to implement a Universal Description, Discovery and Integration protocol (UDDI) web service registry for NOAA. This will allow NOAA developers to register new services and allow developers and the public to search for existing services. Registration information typically provided includes the Web service title, the Web Service description (text), the web service administrator name, e-mail address and phone number, the web service IP address, and access to the web service Web Service Descriptive Language (WSDL) document (which describes the API that the web service exposes to the Internet). The NOAA web service registry will also include information pertinent to the scientific application of the service, for example providing access to a type of data (observation, model, GRID etc.), and the data distribution or accessing a computing resource or application. We also propose to support collaborative efforts by a number of cross-Line Office development projects to develop and register web services.

Implementation

Two types of XML registries are available: Universal Description, Discovery and Integration of Web services (UDDI) or Electronic Business using eXtensible Markup Language (ebXML). This installation will use a web services registry server which implements version 2.0 of the UDDI specification. This implementation will provide a complete solution with regards to a NOAA Web Services Registry. The NOAA Web Services Registry will require the following components and functions:

The registry server: The web service registry server is the core of this application. We will install a public UDDI server that uses version 2.0 of the UDDI specification. Although this server is public, we will only allow for web services from NOAA providers to be registered. The registry will, however, be visible to everyone. Commercial software will be used for this application and we suggest a purchase of the CapeConnect server by CapeClear. This is an enterprise level server, and includes a full public UDDI registry. This server will initially be installed on a Sun Server at PMEL. We will use a NOAA domain (for instance uddi.noaa.gov) for this server that will allow the server to be relocated to any NOAA server. NOAA web service implementers will expose their web service WSDL document to the web service registry during the registration process and thereby provide a full API declaration to the registry.

The registry database: The MySQL database will be used as the storage solution for the CapeClear server.

Service web registration interface: A web interface will be implemented on the CapeClear UDDI server. We will either implement the solution shipped with the server, or modify this to add the functionality we require. The interface will require username and password authentication, and also only allow NOAA agencies to register web services. Additional capability will be done with dynamic PHP pages.

Web access to registry contents: Since human intervention is still pivotal in the discovery of web services, a web interface with search capability of the registry is needed. This web interface will provide on-line functionality to allow consumers of NOAA Web services to browse and search the registry of a NOAA web services. The data access web service search parameters will be by data category (observation, model, gridded, ocean, atmosphere), region (if applicable), line-office (NWS, OAR) and time range. Search results indicate all the information a user needs to connect to the web service, and the information on the API the selected web service exposes. The registry search web interface will also allow a user to search the registry for NOAA web services that provide functionality other than scientific data access (for instance a cloud detection algorithm).

Administration interface: The CapeConnect server ships with a full administrator management interface. This interface is essential to fine-tune the UDDI registry to NOAA's requirements.

Server requirements: The registry will be initially be hosted on a Sun server located on the NOAA Western Region Campus Network, and could be relocated if desired.

Webservice Advocacy

The participants in this proposal believe strongly in the future of web services in NOAA. In order to share knowledge and establish best practices consensus throughout NOAA, members will be involved with partnerships crossing Line Offices to leverage the use of web services and to register and access web services through the UDDI. Several identified partnerships include collaboration between NGDC and SEC, the FSL, PMEL, GFDL grid computing program and a model verification program between NCEP and ETL. These groups will be the initial target users and developers for the UDDI as well as demonstrating the use of web services and UDDI in a number of contexts relevant to NOAA. The following summarizes the development collaborations we propose to support as beta testers of the UDDI.

UDDI application to Grid Computing

FSL has recently become involved in grid computing as a means to utilize compute resources more efficiently, and to simplify access to them. Grid computing relies heavily on web services as a means of resource access, discovery and as a mechanism for data transfer. Access to web services requires knowledge of the data and compute resources. It is not enough to simply make a web service available, it must be discovered by applications and scientists that require them. A centralized data server offers a good way to make these services available to the broader NOAA community. Two FSL proposals have been written to the HPCC program that could use such a server. One proposal, titled "*Development of a prototype NOAA Grid*", proposes to develop a NOAA grid to allow scientists to run general applications that require access to compute and data resources. A second proposal, titled *Real-time Data Cataloging and Distribution*, proposes to build a meta catalogue of data resources in order to allow clients to discover and obtain these data. These proposals would benefit from a centralized UDDI server that would store information about these services and could be queried by scientists when such resources are required.

Near-realtime access to operational data and data fusion

ETL has an existing collaboration with NCEP to provide near-realtime model verification with remote sensor data. This program has required extensive custom data processing to select and package model data. We will assess and prototype a web service to provide access to one of the experimental models.

TAO observational data access

PMEL is in the process of developing a web service that will allow client access to TAO array observational. This web service will be registered and documented in this registry.

Funding has also been requested to bring developers together at a central meeting such as Webshop or NOAAtech where technical presentations can be made to a wider NOAA audience.

Transition to Operations

The Boulder Network Operations Center will participate in this proposal to assess issues related to enterprise architecture, security, and on-going operational deployment. Their assessment will be included in the final report. Continuing costs have been included for 2005 to cover the anticipated transition of the UDDI to a NOC.

Analysis:

Development of a UDDI will help to coordinate web service development efforts, eliminate redundancy and serve as a cornerstone of web services architecture to meet NOAA's growing operational needs.

One alternative to UDDI is Electronic Business using eXtensible Markup Language (ebXML). This implements XML for business applications and specifically targets commercial issues. While the business community and NOAA environments are both data rich and transaction intensive, it is felt that the more generically scoped UDDI will be more easily extended the environmental data and services within NOAA.

We propose to implement the CapeConnect UDDI server, and add a Web interface to allow web service registration and discovery. This enterprise level, commercial server implements a number of standards compliant features and will allow developers to focus on best configuring the service for NOAA usage. The UDDI compliant registry will allow computer clients to register and poll the server, while the web interface will provide a human readable interface to do the same things. Only NOAA web services will be registered on this registry, the public will have access to search the registry. The registry will render all the information to the user that they need to connect to the NOAA web service. Intended for enterprise deployment, this server will meet all of NOAA's current needs for UDDI services. As those needs grow and web services technologies evolve, the standard-based development on this server will transition to the next generation more easily.

Performance Measures:

The proposal will deliver a fully functional industry standard enterprise level UDDI web services registry to serve as a central source of information for web services in NOAA.

Milestones

Month 1 - install and configure server
Month 2 – complete web interface configuration
Months 3-6 – register web services
Months 3-12 - assess UDDI performance and usage
Months 3-12 – conduct outreach to NOAA web services developers
Month 12 – complete operational recommendations

Deliverables:

- UDDI Web Services Registry
- Web interface for web services registration

- Web interface for finding web services in NOAA
- Registration of existing web services.
- Final report with metrics on web service and UDDI usage.