Project Description

Scientific research communities are producing very large databases that are increasingly complex and diverse. These databases are very well suited for particular narrowly-defined, discipline-specific purposes. In principle, they could be used for solving more broadly-defined scientific problems of considerable strategic value, such as understanding ecosystems and human populations. However, repurposing data and knowledge from multiple sources to solve these problems remains a significant challenge.

While databases that are currently being developed typically engender the greatest level of enthusiasm by the communities that are creating them, knowledge discovered in the past can have equal importance for related communities. The problem is not just the differences in formats but also the enormous changes in terminology over time. Databases run the risk of rapid obsolescence as the meaning of the data fields is forgotten even by the individuals who introduced them.

To ensure the strategic value of data and knowledge bases, it is important that they be annotated with metadata. However, annotations alone will not, in themselves, support interoperability between communities and between current terminology and past terminology. It is also necessary to specify high quality, reviewed mappings between terminologies. There is a critical need for a robust, scalable infrastructure for metadata. Such an infrastructure would facilitate interdisciplinary innovation and knowledge discovery.

1 Vision and Rationale

Semantic technologies based on logic, databases and the Semantic Web can address the problem of meaningful access to and integration of data both today and over decades and centuries. This proposal is to develop and deploy a new infrastructure called the Open Ontology Repository (**OOR**) for the registration, storage and management of metadata.[5] This infrastructure will serve a vibrant community of scientific researchers with collections of controlled vocabularies and knowledge models written in XML Schema, RDF, OWL, and other knowledge representation languages.

1.1 Opportunity and Impact

As explained in [18], metadata "summarize data content, context, structure, interrelationships, and provenance (information on history and origin). They add relevance and purpose to data, and they enable the identification of similar data in different data collections." The importance of metadata for scientific research communities was emphasized in a recent joint report of the National Academies of Science and Engineering, and the Institute of Medicine[17] concerned with data usability and integrity, "Enhancing the

Integrity, Accessibility, and Stewardship of Research Data in the Digital Age." This report states, "Metadata are a critical part of the context needed to assess the integrity of data and use data accurately." The report then emphasizes this point when it states, "... and it is generally impossible to judge the integrity of processed data without access to the metadata documenting how they were processed." Moreover, the report states that the role of metadata has become increasingly important over time, "As digital data has become more important in a variety of disciplines and fields, the scope and value of metadata have grown, leading to the development of metadata standards. Metadata standards are an agreed set of terminologies, definitions and values to be provided for data in a given field or community."

Metadata is itself a form of data with its own logical structure and standard formats. Metadata includes both annotations (such as provenance information) and semantics (declarations of classes, properties and constraints). Metadata is organized into collections called *ontologies*, a term that was borrowed from philosophy. While there are important connections between the philosophical notion of ontology and modern ontology engineering, they are distinct fields of study with very different purposes and techniques.

Given the recognition of the importance of metadata by the National Academies, the OOR has the potential for having a major impact on the conduct of modern data-intensive scientific research. However, simply creating the infrastructure and making it available will not automatically ensure that it is utilized in spite of the mandates in the National Academies report cited above. To ensure that the potential of the OOR is fully realized, the project will include a vigorous education and outreach program. The need for expertise in ontology development and application has been recognized by the ontology community. The Ontology Summit 2010[23] was entirely devoted to addressing the need for expanding and improving the training of ontologists.

The OOR project will seek to bring data-intensive research communities into the OOR initiative. Since the OOR will be an open, federated architecture and infrastructure, it is intended to be utilized by communities to host their own ontologies as well as to allow the communities to adapt previously established ontologies for their own purposes. The proposed OOR project will support an initial set of OOR servers, but the intention is that the OOR will consist of a collection of federated servers provided by many communities rather than a single, centralized collection of OOR servers.

The National Academies report[17] raises several concerns about data. One of these is a concern with the long-term preservation of data, "... our focus in this report is on a specific aspect of utility that we refer to as data stewardship – the long-term preservation of data so as to ensure their continued value, sometimes for unanticipated uses. Stewardship goes beyond simply making data accessible. It implies preserving data and metadata so that they can be used by researchers in the same field and in fields other than that of the data's creators." One consequence of this concern is the need to maintain all versions of metadata standards so that preserved data will still be meaningful as standards evolve. Another consequence is the need to map older metadata to newer metadata so that data from different time periods can be compared.

While the OOR would seem like a natural candidate for funding because of its potential impact on scientific research, there are several barriers to overcome.

- Since the OOR would be a general infrastructure which can benefit all data-intensive scientific research communities, it is not within the purview of any particular research community. It would be inappropriate for a single community or small set of communities to provide a service for all communities, possibly including commercial, industrial, government and research communities not funded by the NSF.
- The OOR is an infrastructure for metadata, not data. While it is expected that some data will be supported, it is not the main purpose of the OOR. Image data, for example, would not be appropriate for the OOR. Programs whose focus is on fostering interoperability and collaboration among scientific research communities (such as the OCI INTEROP program) might seem like a good match to the OOR. However, the emphasis of such programs is on data interoperability. Metadata is also crucial, but it is necessarily of secondary importance. The amount of metadata is considerably smaller so that data retrieval, processing and transmission dominate server and network performance.
- As discussed in the first point above, the OOR would be a general infrastructure, so it is difficult to argue that it provides enough specific benefits to a single community to make it competitive with proposals that are specifically targeted at the central problems of a discipline. This is so in spite of the fact that some communities (such as biomedical research communities) are beginning to recognize the need for the OOR and have devoted some resources to its development.

Because of these barriers, the STCI program is the singularly unique opportunity for OOR funding.

The OOR will be developed in a series of phases.

- **Year 1**: Completion of the first production release of the OOR platform, including federation. Initial delivery of OOR services to the targeted communities. Feedback from the communities will be gathered for use in the next phase. Planning begins for the second release. Education and outreach programs expand but remain within the targeted communities.
- **Year 2**: Development of the second release begins, based on lessons learned. Education and outreach programs expand beyond the targeted communities. Delivery of OOR services to additional research communities.
- **Year 3**: Second release of the OOR platform. Education and outreach expands to include commercial and industrial communities.

- **Year 4**: Development of final release that will support all planned features. Education and outreach efforts continue.
- **Year 5**: Final release of the OOR platform with all planned features fully supported. Education and outreach are now the primary roles played by the project.

The OOR team currently meets biweekly using a virtual collaborative environment. This environment is managed by Peter Yim. See Section 2.5. In addition to the biweekly meetings, subteams will have their own meetings, on a weekly or biweekly schedule.

2 Activities

We propose to develop an open ontology repository (OOR) of controlled vocabularies and knowledge models that have been encoded in XML Schema, RDF, OWL, and other knowledge representation languages. The targeted research sector will be the climate and environmental communities.

2.1 Preliminary work

The OOR team has started work on the OOR development. The following are the main activities, each of which is elaborated in more detail in later sections:

- A series of conferences and workshops were held to identify requirements and use cases.[2, 19] A Use Case Description Ontology (UCDO) was developed to encode the results of these meetings.[4] The initial OOR use case descriptions have been published using the UCDO, and development of additional use cases is ongoing to respond to newly identified requirements.[3]
- The initial OOR system architecture was developed based on the use case descriptions. The architecture includes administration, gatekeeping and workflow components along with repository and federation components.
- Several prototype OOR instances have been deployed both at academic institutions such as Northeastern University, University of Bremen and Ryerson University, and at commercial firms including BBN and CIM3.
- Experiments with federation have been conducted at BBN.
- Prototype components for administration, gatekeeping and workflow management have been developed. The administration and gatekeeping components are based on ISO/IEC 11179.

2.2 Metadata life cycle management

To truly support discovery, innovation and learning well into the future, the OOR collaboration will manage the metadata life cycle by providing an architecture and an infrastructure that supports the creation, sharing, searching, and management of ontologies. Complementary goals include fostering the ontology community, the identification and promotion of best practices, and the provision of services relevant to ontologies. Examples of anticipated services include automated semantic interpretation of content expressed in knowledge representation languages, the creation and maintenance of mappings among disparate ontologies and content, and inference over this content. The OOR will support a broad range of semantic services and applications of interest to enterprises and communities.

A series of conferences and workshops were held to solicit input from a large variety of communities. This input was used to identify requirements and use cases. The following were identified as the core requirements for the OOR:

1. The repository architecture shall be scalable.
2. The architecture shall be optimized for sharing, collaboration and reuse.
3. The repository shall be capable of supporting ontologies in multiple formats and levels of formalism.
4. The repository architecture shall support distributed (federated) repositories.
5. The repository architecture shall support explicit machine usable/accessible formal semantics for the meta-model of the repository.
6. The repository shall provide a mechanism to address intellectual property and related legal issues/problems.
7. The repository architecture shall include a core set of services, such as support for adding, searching and mapping across ontologies and data related to the stored ontologies.
8. The repository architecture shall support additional services both directly within the province of the repository and as external services.
9. The repository should support all phases of the ontology lifecycle.

Some of these requirements can be satisfied by using an existing revision control system such as CVS or Subversion. However, there are a number of important requirements of the OOR that distinguish it from generic repositories:

1. Ontologies are the basis for communication between people as well as computer systems. For an ontology to be useful, it must be the result of a consensus among the members of a community. While collaboration tools can be used for developing ontologies, a review process is necessary for achieving community consensus.

Such a process is similar to processes currently in use for developing standards and for publishing research articles.

2. Because the OOR is community-centered, it must allow for variation among communities, not only their ontologies but also the processes they use for the ontology life cycle, in particular for reviewing and publishing their ontologies.
3. To facilitate disciplined reuse, processes must be supported for mapping between ontologies and for composing collections of ontologies to form larger ontologies. Existing repositories have some support for dependencies among artifacts, but the dependencies are generally ad hoc and informal. Ontologies have a formal logical structure that can be employed for formal reasoning. Accordingly, the mappings and compositions must not only be formally specified, they must be administered using the same processes used for the ontologies.
4. The requirement for the OOR to be federated, rather than centralized, has significant consequences for the OOR architecture and design.
5. The knowledge base must have support for inference, consistency checking, composition and differencing.

To address these issues, a collection of actors and use cases was developed. Each use case is fully described as a flow of events involving actors playing appropriate roles. The use case diagram is shown in Figure 1.

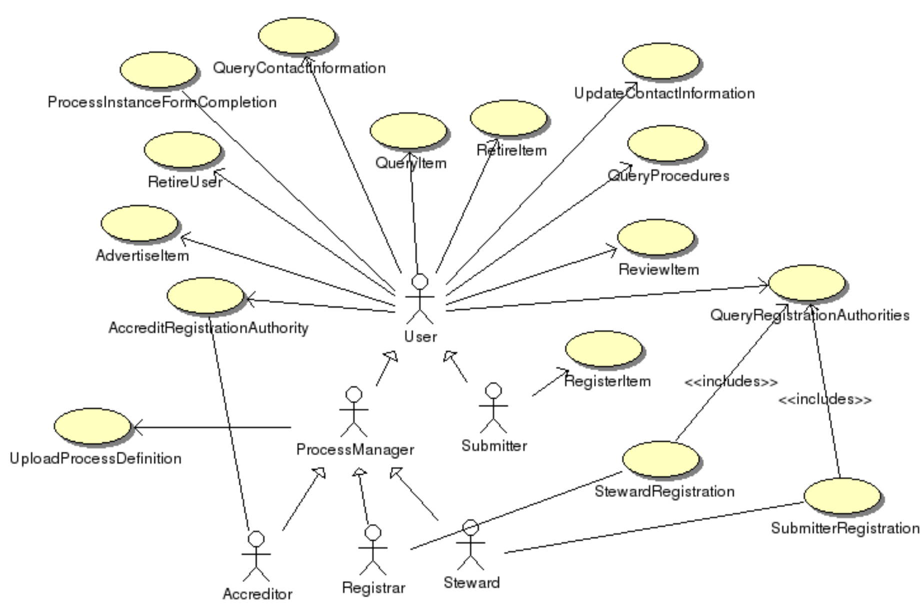

Figure 1: OOR Use Case Diagram

2.2.1 Architecture

The initial system architecture is shown in Figure 2. It is based on the requirements and use cases, especially the OOR-specific issues described in the previous section. The components will be developed as web services supporting both the SOAP and REST web service styles. As a result, each component will have formally specified interfaces. The components can be deployed on multiple machines and replicated as needed to ensure scalability. The proposed project will develop a reference implementation with the hope that other implementations will later be developed.

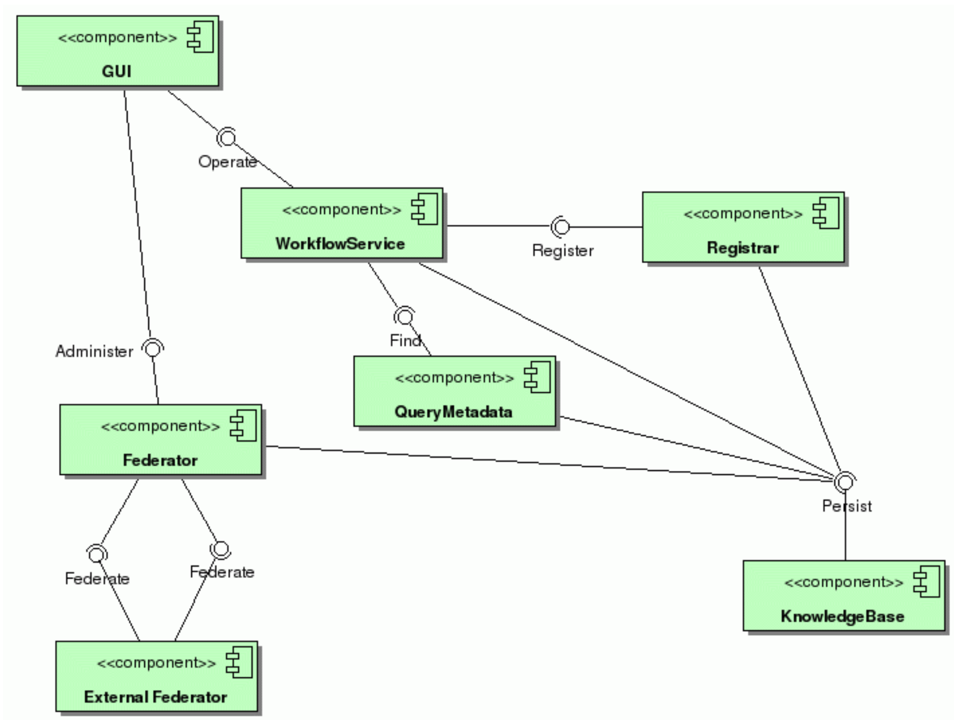

Figure 2: OOR Architecture

The components have the following purposes:

1. **Workflow Service** This is the main entry point for interacting with the OOR. This reflects the fact that non-trivial interactions with the OOR involve multiple steps and can involve more than one actor. Moreover, these processes will vary from one community to another. The Business Process Execution Language (BPEL) will be used to specify these processes.
2. **Query Metadata** The metadata associated with ontologies and other knowledge based artifacts is itself formally defined. This component is responsible for finding appropriate ontologies, mappings and other items in response to specified user needs.

3. **Registrar** The OOR is both a registry and repository. Registration is concerned with the life cycle of ontologies and related items. An item can be submitted, reviewed, harmonized with other items, recommended and possibly retired.
4. **Federator** The OOR is required to be federated. The federator component interacts with the federator component of other OOR instances to ensure consistency of the administered items and the metadata annotations on them.
5. **Knowledge Base** Storage of knowledge based artifacts is complex because of the requirement for supporting inference, consistency checking and other semantic operations. The reference implementation will use existing storage technologies for ontologies (e.g., triple stores).
6. **GUI** Graphical user interfaces will be provided. These will be decoupled from the components providing the main OOR functionality. The reference implementation will adapt existing open-source GUIs.

Metadata occurs on many levels. While it is not always possible or necessary to distinguish these levels, it is expected that the metadata in the repository will usually be organized according to Figure 3. In this figure the data are from 3 different domains. In the first, the data are in the form of image data; in the second, the data are sensor data; and in the third, the data form a knowledge base. The dashed line separates the data being managed by the OOR and data managed separately from the OOR. Both image data and sensor data require domain-specific storage, retrieval and processing software and hardware so they are managed separately. Knowledge bases, on the other hand, are well suited to storage, retrieval and processing by the OOR. Each data set is annotated with metadata that is specific to the data set. This metadata includes provenance information as well as any structural and processing specifications that are specific to the data set. Each domain has its own ontology that specifies those aspects of the metadata annotations that are not covered by the OOR ontology. The domain-specific ontologies and data set-specific metadata are annotated as artifacts in the repository. This is called repository metadata in the figure. The semantics of the repository metadata is specified by the OOR ontology.

The organization shown in Figure 3 is well suited for data sharing and interoperability within one domain, but not for cross-domain interoperability. Achieving the latter requires that the domain ontologies be related to each other. Although the figure does not explicitly show this, the various ontologies will be related to each other in a web of interconnections. Ontology matching and mediation is an active area of research. The OOR will support the specification of mappings between ontologies, both at the level of whole ontologies (e.g., ontology importing) and at lower levels of granularity (e.g., declaring that two resources in different ontologies are the same). The OOR will also provide APIs that allow one to plug in ontology matching modules. However, this project will not be addressing the problem of ontology matching and mediation in an automated or semi-automated manner.

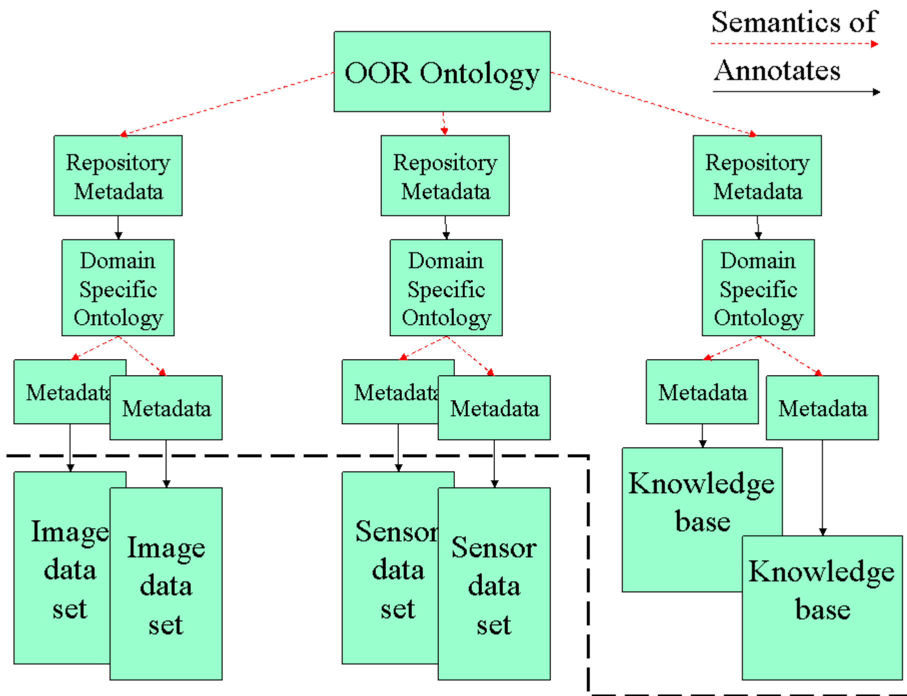

Figure 3: Metadata Levels in the OOR

2.2.2 Ontology interoperability, standards, and integration

To support the sharing and reuse of ontologies within the repository the OOR will store both ontologies and metadata for ontologies. The Ontology Metadata Vocabulary (OMV), Dublin Core, ISO 11179, ISO 19763, and other existing approaches to provenance and versioning support will be used as the basis for the metadata in the OOR. The metadata will allow users to:

- determine whether an ontology is suitable for a user purpose;
- capture the design rationales that underlie the ontology;
- find information about author, author credentials, and source of ontology reference material;
- find reviews and comparisons of ontologies;
- retrieve ontologies for use in domain applications;
- retrieve ontologies to be integrated with other ontologies;
- retrieve ontologies that will be extended to create new ontologies;
- determine whether or not an ontology can be integrated with given ontologies;
- determine whether a set of ontologies retrieved from the repository can be used

together; and

- determine whether an ontology in the repository can be partially shared.

The metadata will include both logical metadata (logical properties of the ontology independent of any implementation or engineering artifact) and engineering metadata (properties of the ontology considered as an engineering artifact).

2.3 Research Problems

Ontology engineering is a rapidly developing research area. There are both computational and storage challenges as the size of a knowledge base grows, since the computational complexity of logical inference is much greater than the complexity of traditional database query processing. The OOR team is closely connected both with the ontology engineering community and with the communities specifically targeted by this proposal. As shown in the biosketches, the team members have substantial ontology research experience. Ontology engineering research will be an integral part of the OOR effort. In particular, Mike Dean, an OOR co-convener, will engage in research on federation among OOR and non-OOR registries, repositories, and collaborative development environments (including the emerging area of Semantic Wikis).

Ontology repositories constitute a subfield within ontology engineering, and has its own community and conferences. In a recent meeting chaired by Ken Baclawski, the state of the art of ontology repositories was evaluated and the research challenges were identified[1]. The proposed project will focus on the issues of federation, modularization and education. These were selected because they are the ones that are most likely to increase the level of adoption of the OOR by scientific research communities. They are also the issues that have not yet been adequately addressed by existing ontology repository projects.

2.4 Education and Training

As an integral part of the proposed project, the OOR will support a vigorous educational outreach program to bring data-intensive research communities into the OOR initiative. Since the OOR will be an open, federated architecture and infrastructure, it is intended to be utilized by communities to host their own ontologies as well as allowing the communities to adapt previously established ontologies for their own purposes.

The Ontolog Forum has been engaging in educational and outreach activities for 7 years, reaching over 40 distinct communities. Examples include communities in bioinformatics, national command and control, and intelligence [16, 15, 14, 21, 20, 24, 26, 27, 28, 29].

2.5 International community and user input and assessment

This OOR initiative has emerged from a collaborative effort by a community that has successfully developed an infrastructure which has grown and sustained itself since early 2002. Ontolog (or the Ontolog Forum) is an open virtual community for ontology engineering, with members who realize the importance and potential impact of ontology, and are passionate about moving it into the mainstream through adoption and standardization. With over 550 participating members from around 30 different countries, Ontolog has been able to amass a highly regarded body of knowledge in the ontology domain. This was done by archiving (and metadata tagging) all its membership dialog and contributions transacted over its collaborative work environment infrastructure, which also serves as a dynamic knowledge repository for the community's collective intelligence. This open repository receives about 50,000 page views or file downloads per day, from around 120 cities worldwide [30].

The "Open Ontology Repository" theme was selected by the community as the theme for their 2008 Ontology Summit. This is an annual event (since 2006) jointly organized by NIST, Ontolog and the National Center for Ontological Research. The Ontology Summit entailed a three-month online discourse, four virtual panel discussions, a two-day face-to-face workshop at NIST, and the publication of a Communiqué. The initiative was co-sponsored by about 50 institutions, and had an Organizing Committee and an Advisory Committee made up of more than 40 individuals, who are among the most respected names, worldwide, in the ontology and ontological engineering domain [22].

3 Organizational Structure

3.1 Leadership and Management

Kenneth Baclawski will be responsible for overall coordination of the project and for directing the outreach to scientific research communities. Mike Dean will be the PI for the BBN subcontract that will provide open source software modules to support federation of ontology repositories. Peter Yim will be the Operations Director for the OOR servers to be supported by the project. Peter brings along an illustrious visionary, strategic leadership and cross-disciplinary management portfolio. He has managed as many as 3000 people and 300 software engineers, and has built companies from scratch to $500 million in revenue. He has learned the art and science of working effectively at almost any part of a real or virtual organization – from the executive office, to corporate boardrooms, to R&D, all the way to the shop floor. He works in corporate settings and academia, as well as on non-profits, education and government boards. He has been a co-founder of Ontolog Forum (along with Leo Obrst); the OASIS Universal Business Language (UBL) standard technical committee; Director of Program Management for VerticalNet, Inc., the first B2B company that employed ontologies for data integration in an eCommerce setting; as well as having been a principal or chief executive at other industrial enterprises. He now heads CIM3, a distributed collaboration technology and internet service company which has been providing infrastructure support to the Ontolog Forum operations, as well as to various US Federal inter-governmental collaborative endeavors.

3.2 Comprehensive expertise and infrastructure capacity/capabilities

The OOR team possesses considerable expertise in library and archival sciences; computer, computational and information sciences; and cyberinfrastructure. In addition, individual team members have expertise in particular domains as discussed in Section 3.3 below. The team has been collaborating at virtual meetings for over 2 years, as well as at a face-to-face meeting at NIST in April of 2008. The team has developed an effective organizational structure that enables shared responsibility, close coordination and cooperation, and catalyzes the rapid exchange of ideas. The team has access to considerable computational, storage, network access, dissemination, interaction and communication resources. Peter Yim is responsible for managing this infrastructure as explained in the previous section above.

3.3 Diverse, multi-sector participation

The initial sectors that will be served by the OOR include the climate and environmental communities, with a focus on marine environments, as well as the biomedical communities that are already served by BioPortal. BioPortal primarily serves the biomedical research sector, including the genomics, proteomics, diseases, anatomy, and model organism communities. The OOR team is currently developing a prototype of the server based on BioPortal, which is a centralized repository for biomedical ontologies.[19, 25, 8]

The Marine Metadata Interoperability (MMI) project team will collaborate with the OOR project to integrate MMI project work with the OOR effort. MMI has as its key mission objective developing a broad community presence to address marine metadata issues. The MMI project is uniquely positioned to understand community needs, address those needs in open and broadly applicable ways, and obtain community engagement and buy-in for open solutions[7, 6]. With Technical Lead Carlos Rueda (Monterey Bay Aquarium Research Institute) and contributor Luis Bermudez (Southeastern Universities Research Association), among many others, MMI's pioneering work developing a community semantic architecture and ontology repository will provide key insights and building blocks for the OOR project[9, 10, 11, 12]. MMI's marine science solutions are equally valuable and viable in most environmental communities, and a natural progression of the effort; several environmental science communities are represented at their semantic interoperability workshops[13].

Ultimately the OOR initiative will be providing:

- Guidance for data providers, metadata providers, and ontology providers;
- Organized references on all facets of metadata needs and solutions;
- Services targeting semantic interoperability in the respective and related domains, including vocabulary lists, ontology repository and associated services, and vocabulary creation and maintenance tools, services, and guidance;
- Community collaboration environment (shared files, email archives, and web pages, either public or secure);
- Access to work in progress on metadata tasks and projects.

In addition, the OOR collaboration can provide purposeful capabilities:

- Targeted identification and evaluation of resources (vocabularies, standards, tools, services);
- Identification and engagement of key community participants (projects or individuals) in metadata initiatives;
- Training and workshops in metadata technologies and techniques, particularly dealing with semantic tools and services, including vocabulary and ontology development, metadata standards and their application, as well as metadata-enlightened architectural development and analysis;
- Community environment(s) to advance particular topics or discussions.

The references section of the proposal includes links to resources that are already available in the targeted sector.

References

[1] K. Baclawski. Ontology repository research issues workshop, August 2009. http://ontolog.cim3.net/cgi-bin/wiki.pl?ConferenceCall_2009_08_06.

[2] K. Baclawski. Ontology repository use case workshop, April 1 2010. http://ontolog.cim3.net/cgi-bin/wiki.pl?ConferenceCall_2010_04_01.

[3] K. Baclawski. Open ontology repository use case descriptions, 2010. http://www.ccs.neu.edu/home/kenb/ontologies/oor-usecase.xml.

[4] K. Baclawski. Use case description ontology, 2010. http://www.ccs.neu.edu/home/kenb/ontologies/ucdo.owl.

[5] K. Baclawski and T. Schneider. The open ontology repository initiative: Requirements and research challenges. In *Collaborative Construction, Management and Linking of Structured Knowledge*, October 2009.

[6] L. Bermudez, J. Graybeal, and R. Arko. A marine platforms ontology: Experiences and lessons. In *Proceedings of the Semantic Sensor Networks Workshop at the 5th International Semantic Web Conference ISWC 2006*, Athens, GA, November 5-9 2006.

[7] L. Bermudez, J. Graybeal, A. Isenor, R. Lowry, and D. Wright. Construction of marine vocabularies in the marine metadata interoperability project. In *Proceedings of the 13th Ocean Sciences Meeting*, Washington, D.C., September 19-23 2005.

[8] M. Dean. Coordinating our open ontology repository software development, February 19 2010. http://ontolog.cim3.net/cgi-bin/wiki.pl?OOR/ConferenceCall_2010_02_19.

[9] J. Graybeal. Marine Metadata Interoperability Community Vocabulary Use Cases, 2008. http://marinemetadata.org/usecasesvocs2008.

[10] J. Graybeal. Marine Metadata Interoperability Ontology Registry, 2008. http://marinemetadata.org/community/teams/ont/mmirepository.

[11] J. Graybeal. Marine Metadata Interoperability Ontology Registry and Repository, 2008. http://mmisw.org/or.

[12] J. Graybeal. Marine Metadata Interoperability Semantic Framework, 2008. http://marinemetadata.org/semanticframeworkconcept.

[13] J. Graybeal. Marine Metadata Interoperability Semantic Interoperability Workshop Series, 2008. http://marinemetadata.org/workshops/oosinterop2008.

[14] J. Luciano. A pathway to pathways. In *Genome Technology*, pages 15–16, June 2006.

[15] J. Luciano and R. Stevens. e-Science and biological pathway semantics. *BMC Bioinformatics*, 8(Suppl 3), 2007.

[16] J. Luciano and R. Stevens. OWL: PAX of mind or the AX? Experiences of using OWL in the development of BioPAX. In *OWL: Experiences and Directions*, Gaithersburg, MD, USA, April 1-2 2008.

[17] *Enhancing the Integrity, Accessibility, and Stewardship of Research Data in the Digital Age*. National Academies Press, July 22 2009.

[18] Cyberinfrastructure Vision for the 21st Century, 2007. `http://www.nsf.gov/pubs/2007/nsf0728/index.jsp`.

[19] N. Noy. BioPortal and related ontology repository implementations and development, July 30 2009. `http://ontolog.cim3.net/cgi-bin/wiki.pl?OOR/ConferenceCall_2010_02_19`.

[20] L. Obrst, T. Janssen, and W. Ceusters, editors. *Ontologies for Intelligence Analysis*. IOS Press, 2008. forthcoming.

[21] L. Obrst, D. McCandless, S. Stoutenburg, K. Fox, D. Nichols, M. Prausa, and R. Sward. Evolving use of distributed semantics to achieve net-centricity. In *Regarding the "Intelligence" in Distributed Intelligent Systems*, Arlington, VA, Nov. 8-11 2007.

[22] April 2008. `http://ontolog.cim3.net/cgi-bin/wiki.pl?OntologySummit2008`.

[23] March 2010. `http://ontolog.cim3.net/cgi-bin/wiki.pl?OntologySummit2010`.

[24] M. Parmelee, D. Nichols, and L. Obrst. A net-centric metadata framework for service oriented environments. *International Journal of Metadata, Semantics and Ontologies (IJMSO)*, 2008. forthcoming.

[25] T. Schneider. OOR-BioPortal comparative analysis, October 15 2009. `http://ontolog.cim3.net/cgi-bin/wiki.pl?ConferenceCall_2009_10_15`.

[26] S. Semy, M. Pulvermacher, and L. Obrst. Toward the use of an upper ontology for U.S. government and U.S. military domains: An evaluation. Technical Report MTR 04B0000063, MITRE, November 2005. `http://www.mitre.org/work/tech_papers/tech_papers_05/04_1175/index.html`.

[27] S. Stoutenburg, L. Obrst, D. McCandless, D. Nichols, P. Franklin, M. Prausa, and R. Sward. Ontologies for rapid integration of heterogeneous data for command, control, and intelligence. In *Ontologies for the Intelligence Community Conference*, Columbia, MD, Nov. 28-30 2007.

[28] S. Stoutenburg, L. Obrst, D. Nichols, P. Franklin, K. Samuel, and M. Prausa. Ontologies and rules for rapid enterprise integration and event aggregation. In *Vocabularies, Ontologies and Rules for the Enterprise (VORTE 07), EDOC 2007*, Annapolis, MD, Oct. 15-19 2007.

[29] S. Stoutenburg, L. Obrst, D. Nichols, K. Samuel, and P. Franklin. Applying semantic rules to achieve dynamic service oriented architectures. In *RuleML 2006: Rules and Rule Markup Languages for the Semantic Web*, volume 4294, pages 581–590, Heidelberg, November 10-11 2006. Springer-Verlag. http://www.mitre.org/work/tech_papers/tech_papers_06/06_0904/index.html.

[30] P. Yim et al. Ontolog's collaborative work environment and knowledge repository, 2002–2010. http://ontolog.cim3.net/wiki.