

Notes from the Interoperability Front: A Progress Report on the Open Archives Initiative

Herbert Van de Sompel¹, Carl Lagoze²

¹Research Library, Los Alamos National Laboratory,
Los Alamos, NM
<mailto:herbertv@lanl.gov>

²Computing and Information Science, Cornell University
Ithaca, NY USA 14850
lagoze@cs.cornell.edu

Abstract. The Open Archives Initiative Protocol for Metadata Harvesting (OAI-PMH) was first released in January 2001. Since that time, the protocol has been adopted by a broad community and become the focus of a number of research and implementation projects. We describe the various activities building on the OAI-PMH since its first release. We then describe the activities and decisions leading up to the release of a stable Version 2 of the OAI-PMH. Finally, we describe the key features of OAI-PMH Version 2.

1 Introduction

Over a year has passed since the first release of the Open Archives Initiative Protocol for Metadata harvesting (OAI-PMH) in January 2001. During that period, the OAI-PMH has emerged as a practical foundation for digital library interoperability. The OAI-PMH supports interoperability via a relatively simple two-party model. At one end, *data providers* employ the OAI-PMH to expose structured data, *metadata*, in various forms. At the other end, *service providers* use the OAI-PMH to harvest the metadata from data providers and then subsequently automatically process it and add value in the form of services. While resource discovery is often mentioned as the exemplar service, other service possibilities include longevity and risk management [19], personalization [16], and current awareness.

The general acceptance of the OAI-PMH is based on a number of factors. It is intentionally low-barrier, exploiting widely deployed Web technologies such as HTTP and XML. It builds on many years of metadata practice, leveraging the development of a lingua franca metadata vocabulary in the Dublin Core Metadata Initiative [4]. It accommodates a number of community and domain-specific extensions such as the co-existence of multiple domain-specific metadata vocabularies, collection descriptions, and resource organization schemes.

The first sixteen months after the Version 1 release of the OAI-PMH were purposefully experimental. The intention during that time was to provide a reasonably stable platform for early adopters to test the concepts of metadata harvesting and build a number of fundamental services. Indeed, that intended

stability was accomplished, with only one change in the protocol occurring over the sixteen months due to a change in the XML Schema specification [17].

This paper reports on the results of that approximately year and a half of experimentation and the follow-on activity leading up to a stable Version 2 release of the OAI-PMH. This paper serves as a direct follow-on to an earlier paper [21] that provides background information on the initial development of the OAI-PMH. That historical material is not repeated in this paper and readers new to the OAI are invited to first read that paper.

This paper is organized as follows. The next section of the paper, Section 2, describes the communities of implementers, service providers, and researchers that have developed since the release of the OAI-PMH. Section 3 then describes the process of developing Version 2 of the OAI-PMH and enumerates the changes incorporated into that new version. The paper closes with Section 4 that describes possible next steps and the general future of the OAI.

2 Measures of success

Measuring the success of the OAI-PMH is problematic. Unlike user-oriented technologies (e.g. word processors, spreadsheets), there are no immediate or direct benefits from individual adoptions. Therefore, while a simple count of protocol implementations is indeed useful as a metric, other factors should be considered.

Using terms presented by Shapiro and Varian [25], technologies such as the OAI-PMH exhibit *network effects*, in which initial adoption may be slow and steady and positive feedback then dramatically increases the adoption rate. In the context of the OAI-PMH, positive feedback occurs at two levels. First, it comes from the establishment of communities and/or research projects that collectively agree to adopt the protocol as a basis for information federation. Second, positive feedback comes from the growth of service providers who encourage data provider implementations by adding value to the metadata that they provide.

The past sixteen months indicate that the bases of this positive feedback loop are indeed being established. The remainder of this section describes the activities that contribute to this. First, we give a summary of some notable research grants and projects that provide a context for individual implementations. We then describe a number of service provider implementations. In the next section, we describe a number of the tools that make it easier for, and thus encourage, data providers to undertake repository implementations. We close with a summary of the growth of data providers that has been encouraged by these other developments.

2.1 Funded Research Projects and Programs

Over the past year a number of US and European research initiatives and projects have been established that apply the metadata harvesting model provided by the OAI-PMH. Collectively these projects provide a strong incentive for data providers to reveal their metadata via the OAI-PMH. The following list includes some of the more notable ones.

National Science Digital Library (NSDL). The NSDL is a National Science Foundation funded project to build what will probably be the largest and most diverse digital library to date. Over the next five years the library is expected to serve millions of users and provide access to tens of millions of digital resources. Complete details on the NSDL are provided in a set of earlier papers [10, 20, 31].

The OAI-PMH plays a fundamental role in the NSDL architecture, and the expected size of the NSDL will make it one of the primary deployment contexts for the protocol. A core component of the NSDL architecture is a Metadata Repository (MR) that provides robust, central storage of metadata in multiple formats related to resources and collections in the NSDL. Figure 1 illustrates the role of OAI-PMH in relation to the NSDL MR. On the ingest side (the bottom of the figure) OAI-PMH is the preferred mechanism for retrieval of collection and resource level metadata from participating data providers. A “front porch” then normalizes and crosswalks (to other preferred formats) this metadata and metadata collected via other means (gathering and direct entry). The processed metadata is then transferred via OAI-PMH to the MR. Finally, the MR exposes its multiple format metadata via OAI-PMH for harvesting by services that provide the bulk of NSDL functionality (e.g., search and discovery, preservation).

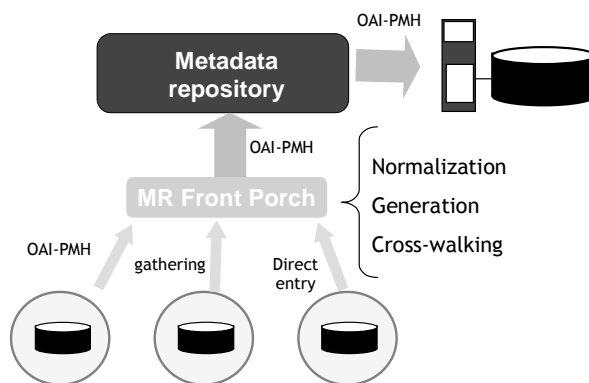


Fig. 1 Role of OAI-PMH in the NSDL Metadata Repository

The first production release of the NSDL is scheduled for 3rd quarter 2002 with subsequent releases over the next four years. We expect that this will be reflected in a dramatic increase in the number of OAI-PMH compliant data providers.

Metadata Harvesting Initiative of the Mellon Foundation. In August, 2001 the Andrew W. Mellon Foundation awarded seven grants totaling \$1.5M to fund development of services on top of the OAI-PMH infrastructure [29]. The awarded projects address a variety of scenarios including:

- designing portal services based on metadata from multi-disciplinary and multi-institutional domains.
- harvesting metadata from archives and special collections.
- harvesting metadata related to materials in a specific topic area, but which are in a variety of formats.

The awardees are Research Libraries Group, University of Michigan, University of Illinois at Urbana-Champaign, Emory University, Woodrow Wilson International Center for Scholars, University of Virginia, and Southeastern Library Network. At the time of writing this paper (May, 2002), many of these projects are still in the start-up phase although preliminary results are available from some [6, 15]. We expect that these projects will help develop guidelines for future service providers and help clarify the relationship between metadata quality and service functionality.

Open Language Archives Community (OLAC). OLAC is a distributed, federated archive of language resources [12]. There are currently fourteen repositories that support an extended version of the Dublin Core metadata set and a compatible but profiled version of the OAI-PMH [11]. While interesting from a technical standpoint, the most intriguing aspect of OLAC is the rich scholarly community that the technical infrastructure has engendered. In a real sense this is what infrastructure like the OAI-PMH is about – not as an object of attention in itself, but as an almost invisible catalyst for far more interesting activities.

eprints.org. The eprints.org self-archiving software has been developed by the Electronics and Computer Science Department at the University of Southampton (UK) [5]. The software has been built as basic infrastructure in support of a general model of “author self-archiving”. This model is proposed by Steven Harnad [18] as a means of reforming the scholarly publishing framework. eprints is available for free under an open source license to individuals, institutions, and learned societies who wish to set up an archive for submission, storage, and dissemination of scholarly publications. The software is fully OAI-PMH-conformant, thus enabling the open federation of scholarly eprint archives that inspired the origination of the OAI [27, 28]. eprints is currently installed at approximately thirty institutions and the recent release of Version 2 of the software (February 2002), featuring significantly easier installation, promises to ensure the rapid growth of these OAI-PMH-conformant repositories.

European OIA-PMH Funding Programs. During the past year a number of programs providing funding for OAI implementations and services have been announced in Europe. These include:

- *JISC FAIR* - In January, 2002 the Joint Information Systems Committee in the UK announced the Focus on Access to Institutional Resources Programme (FAIR) [15]. This program will provide considerable funding for projects to explore dissemination of institutional assets or create services via the OAI-PMH.
- *DINI* - Die Deutsche Initiative für NetzwerkInformation (DINI) (German) is funding projects that implement and use the OAI-PMH [6].

- *Open Archives Forum* – OAF is an EU-funded accompanying measure involving the University of Bath – UKOLN (UK), Istituto di Elaborazione della Informazione-CNR (Italy) and Computing Center of Humboldt University (Germany) [7]. The goal of the project is to facilitate the creation of an OAI community in Europe through workshops and support activities. The first workshop of the OAF is being held in Pisa in July, 2002.

Museum Community Projects. The CIMI Consortium [3] has been working with members to support the deployment of the OAI-PMH in the museum community [24]. As part of this process CIMI has developed a number of OAI-PMH tools for the museum community and run workshops at major museum conferences. At present, the most active area of deployment is in Australia led by Australian Museums OnLine, which is building a search and discovery system over the collections of member museums [2].

2.2 OAI-PMH-related Tools

A key factor in the growth of implementations of the OAI-PMH has been the availability of tools to build OAI-PMH-conformant repositories and harvesters. In this section, we enumerate some of the available toolsets.

Repository Explorer. The Repository Explorer [26] is an easy-to-use interactive tool that allows repositories to test conformance of their implementation of the OAI-PMH. The interface available on the Web at http://purl.org/net/oai_explorer allows implementers to enter the base-URL of their OAI repository and then test responses to each verb with varying arguments. The availability of the Repository Explorer has proven to be extremely valuable for new implementers.

ALCME. The Advanced Library Collection Management Environment (ALCME) is a set of OAI-related tools developed by OCLC. The tools include OAICAT, an open source metadata server that can be placed on top of existing databases to turn them into OAI repositories with minimal coding effort, an OAI harvester, and a MARC to DC translator. The toolset is available at <http://alcme.oclc.org/index.html>.

OAIB. The Open Archives In a Box is an application for exporting metadata from a relational database system via the OAI-PMH. OAIB was developed as part of the University of Illinois at Urbana-Champaign Metadata Harvesting Project. The toolset is available at http://emerge.ncsa.uiuc.edu/documentation_oaib.html.

DP9. This is a tool that allows exposure of the metadata in OAI-PMH-conformant repositories to conventional web-search engines [23]. The concept behind DP9 is deceptively simple but ingenious. The data provider exposes metadata via OAI-PMH, DP9 harvests it and creates a static HTML page with hard links, which actually are OAI-PMH GetRecord requests, to metadata records in the harvested repository. Web-crawlers then encounter this page in normal crawling activity, follow the links provided in the page and as such obtain the metadata from the repository for indexing. As a significant portion of the Web is invisible to search engines due to robots.txt exclusions or the lack of fixed URLs for resources that are dynamically served via database queries [22], DP9 allows a content provider hidden in such a manner to make metadata about content visible to search engines (e.g., Google), and thus searchable by users. DP9 description and tools are available at <http://arc.cs.odu.edu:8080/dp9/about.jsp>.

2.3 Service Providers

In addition to the funded projects described earlier, there have been a number of individual efforts to develop services that process metadata harvested via the OAI-PMH.

SCIRUS. This is an Internet search tool, developed by Elsevier Science that focuses on scientific content. Its corpus includes both web and journal content. The SCIRUS engine also indexes metadata of all 180,000 articles in arxiv.org which is harvested via OAI-PMH. Harvesting other OAI-PMH-conformant repositories is planned.

CDL eScholarship Repository. This a project by the California Digital Library to provide faculty with a mechanism to support pre-publication scholarship. The repository supports OAI-PMH and thus can be federated with similar repositories at partner institutions.

my.OAI. This a search engine that harvests metadata from selected repositories and offers the user options to personalize the service and create alerting profiles. my.OAI is at <http://www.myOAI.com>.

2.4 Data Providers

Beginning the first release of the protocol in January 2001, the OAI has maintained a registration service both for tracking implementations and as a convenience for data and service providers. Registration is not required (there would be no way to enforce it anyway!). Sites that request registration must first pass a conformance test. All sites that successfully complete this test are then listed in browsable form at <http://www.openarchives.org/Register/BrowseSites.pl> and in an XML formatted list at <http://www.openarchives.org/Register/ListFriends.pl>. This XML list is used by a number of the service providers described above.

Figure 2 shows the steady growth of registered data providers since the introduction of the protocol. Historical statistics on the growth of the number of metadata records harvestable via OAI-PMH are not available. However, the current (May 2002) count is approximately 6 million.

It is obvious that the “network effects” described earlier have not yet occurred, and there is no evidence to predict it will. We note two issues. First, a number of service provider projects have just recently gotten up-to-speed or have not yet reached implementation level (the NSDL project is a notable example). Second, the protocol is actively being used in environments where metadata is being shared amongst a restricted amount of nodes, and where public advertisement of the metadata collection is considered inappropriate. Anecdotal evidence indicates that the number of registered sites represent less than half of the actual implementers of the OAI-PMH.




Fig. 2 Growth of data providers in the OAI registry

3 OAI Protocol for Metadata Harvesting Version 2

As noted earlier, Version 1.x¹ of the protocol was explicitly experimental. Results over the past sixteen months, described in Section 2, indicate that the underlying technical scope of the OAI-PMH was properly defined. As described in [21], this scope was motivated by the belief that widespread adoption of the protocol (and any technical infrastructure) depends on ease of understanding and implementation. As such the original protocol specification left controversial issues such as access management out-of-scope and when possible chose simple approaches rather than more functional, but complex approaches.

The decision process for Version 2.0 of the protocol (described in Section 3.1), and the results of that decision process (described in Section 3.2), reflect the results of that experimental period. Major functional changes were determined to be out-of-scope and the focus was directed toward fine-tuning existing functionality and improving the general clarity and overall consistency of the protocol document.

3.1 Protocol Review and Definition Process

An important information tool during the length of the Version 1.x experimental process was the OAI-implementers [1] list that was activated at the time of the initial release. The list and its archive quickly became a valuable source of information for newcomers to the protocol, as well an active forum for detailed discussion of protocol related issues by expert implementers. Those discussions had a considerable influence on the process of reviewing the protocol uncovering missing functionality or parts of the protocol that were confusing.

The first step in the process of officially defining Version 2 of the OAI-PMH was the creation, in June 2001, of a new OAI Technical Committee², OAI-tech. The Committee is a group of expert implementers many of whom had been involved in the creation and/or alpha testing of v.1.0 of the OAI-PMH. The group has twelve representatives from US institutions and four from European institutions. The charge for OAI-tech was:

- To review the details of v.1.x of the protocol and determine whether its functionality and nature was correct;
- To compile a stabilized version of the Metadata Harvesting Protocol by mid-2002.

The entire OAI-tech process has been conducted online or via telephone. It began in September 2001, when OAI-tech members were invited to submit brief descriptions of issues they felt required discussion. These issues were motivated by their own implementation experience and from discussions in the OAI-implementers list. The

¹ Version 1.0 was released in January, 2001. A subsequent update release Version 1.1, required due to changes in the W3C XML Schema specification, was released in June 2001.

² An earlier incarnation of the OAI technical committee was instrumental in the creation of OAI-PMH Version 1.0.

initial issues list was then discussed in a conference call during which advocates for each issue, willing to submit a white paper, were solicited. Those issues for which no one volunteered to write a white paper about were dropped. Then followed a process of detailed discussion per white paper, and a proposal for resolution by the OAI Executive³, which itself was again commented upon by the group.

By February 2002, most issues had been resolved through this process. A few issues remained unresolved and these were taken offline for review and resolution by the OAI Executive who worked with an appointed writing sub-committee of OAI-tech – Michael Nelson and Simeon Warner – to draft a first alpha version of the revised protocol. That version was released to OAI-tech early March 1st 2002, at which point the group was extended with a relatively diverse group alpha-testers [8]. In-depth on-list discussion of the protocol document was conducted throughout March and April 2002, which led to a number of significant changes both in text and in function⁴. On May 1st 2002, the OAI-PMH 2.0beta was released to OAI-implementers. On June 1st 2002, Version 2.0 was officially released, 1 month behind schedule.

3.2 Features and Changes in OAI-PMH Version 2

Clarifications and Specification Cleanups. In general the major changes introduced in Version 2 of the protocol are of the nature of clarifications to ambiguities or better means of expressing existing functionality. The following list summarizes these changes.

Core Protocol versus Extensible Notions. A clear boundary has been established in Version 2 between the core OAI-PMH protocol and peripheral notions. This has been achieved by the creation of Implementation Guidelines documents, which covers issues such as the XML Schema for domain-specific profiling of the OAI-PMH through the usage of container structures provided by the core protocol. Such information had previously been part of the protocol document.

Clarified Data Model. Version 2 of the protocol builds on a well-defined “resource – item – metadata-record” data model, correcting a level of ambiguity that existed with this respect in v.1.x. This correction is a consequence of extending the applicability of the OAI-PMH beyond the exchange of metadata about document-like objects, for which the previously established notions about the nature of full content in repositories were meaningful. It is also a consequence of the increasing prevalence of metadata aggregators, such as the NSDL Metadata Repository described in Section 2.1, that have a distant relationship to actual full content.

Metadata in the OAI-PMH is now about resources in general. As a result an item is now defined as a constituent of a repository from which metadata about a resource

³ The OAI-Executive is Herbert Van de Sompel and Carl Lagoze (the authors of this paper).

⁴ One example of a protocol notion requiring extensive review during the alpha process related to idempotency of resumption tokens, described later in this document.

can be disseminated. The notion of an item now plays an important role in the OAI-PMH, as it has become the logical point of entry to physical metadata records.

Error Reporting. Version 2 establishes a clear separation between errors occurring at the HTTP transport layer used by the OAI-PMH and errors at the OAI-PMH level. This is achieved by introducing a set of OAI-PMH error and exception conditions, which can be reported in a dedicated element in the XML responses to OAI-PMH requests.

Dublin Core Metadata Definition. In a successful and trend-setting collaboration with the Dublin Core Metadata Initiative (DCMI) [4], an XML Schema for unqualified Dublin Core has been created, which is hosted by the DCMI and used in the delivery of metadata in the mandatory DC format in the OAI-PMH. This Schema replaces the one previously created and maintained by the OAI.

Single XML Schema for OAI-PMH Responses. OAI-PMH now defines a single XML Schema to validate responses to all OAI-PMH requests. This single schema has a number of advantages relative to the schema per verb strategy in Version 1: it eliminates redundancy in type definitions, allows for clean handling of OAI-PMH errors, and simplifies protocol implementations.

Language Cleanup. The OAI-PMH now uses the notions *must*, *must not*, *may*, *should*, etc. as in RFC2119 [14], allowing for a better understanding of protocol conformance.

Error Corrections. In addition to the clarifications listed above, a number of other changes are introduced in OAI-PMH v.2.0 that correct flaws exposed during the period of experimentation with v.1.x.

Uniform Date and Time Encoding. Dates and times are now uniformly encoded using ISO8601 [30] and are expressed in UTC throughout the protocol. When time is included, the special UTC designator ("Z") must be used. This change corrects the fact that v.1.x required repositories to express the date/time of responses in *local* time, encoded in ISO8601. However, the time zone of the timestamp, which is crucial for harvesting, was left undefined. This caused synchronization problems between harvesters and repositories.

Flow Control Idempotency. The experimentation period revealed a number of problems harvesting large result sets in the face of transaction failures. In the OAI-PMH, large lists are typically broken up into several incomplete lists. At the end of each incomplete list, a harvester receives a `resumptionToken`, which it can issue to receive the next incomplete list. When, as a result of some error, a harvester does not receive a response to a request with a `resumptionToken`, v.1.x does not provide any guarantees that re-issuing the request with the `resumptionToken` returns the response

that had failed. In order to fix this, Version 2 clearly defines and mandates the idempotent nature of resumptionTokens, providing deterministic behavior in the face of transmission errors.

New Functionality. In a number of areas new functionality has been introduced in OAI-PMH Version 2.

Multiple Time Granularity. In response to requests originating mainly from the union-catalogue realm, the OAI-PMH now supports harvesting at different levels of time-granularity: support of year, month and day granularities is mandatory, while support for hour, minutes and seconds granularity is optional.

Enhanced Identify Semantics. The response to the Identify request is more expressive than in Version 1.x, providing more information to harvesters and supporting multiple harvesting strategies. The response now includes information on the earliest datestamp used in a repository, HTTP-level compression schemes supported by a repository, an indication of the finest granularity with which a repository can be harvested, and information about the support that the repository has for deleted items.

Item Set Membership. For repositories that support set-structures, it is now mandatory to list set membership of items in the responses to the GetRecord, ListRecords and ListIdentifiers requests. This allows harvesters to gather set-membership information for an item in a single request.

Related Changes and Activities. Version 2 introduces a number of interesting additions or changes ancillary to the core protocol.

Implementation Guidelines. Implementation Guidelines documents are provided to support interpretation of the protocol document by implementers of repositories, harvesters, aggregators, mirrors, caches, etc.

Provenance Schema. Version 2 provides a recommended provenance XML Schema that is useful in situations where harvested records are aggregated and subsequently made available for re-harvesting. When re-exposing a record, a provenance container can be attached to it, which unambiguously identifies the origin record. The provenance schema allows listing multiple origins for a given record. The schema is designed to be useful for de-duplicating harvested datasets.

Friends Schema. A “friends” XML Schema is recommended to help establish a dynamic approach to the process of discovering repositories. In the response to the Identify request, repositories can use the friends schema to list the BASE-URLs of other repositories they know about. If successful, this approach would allow harvesters to assemble a comprehensive list of harvestable targets by jumping from repository to repository. This may provide a sustainable alternative to a central

registration service as currently operated by the OAI.

Mini-Archive Specification. Version 2 provides an XML Schema that defines an XML document format containing all the information required to respond to OAI-PMH requests. The so-called mini-archives specification comes with an accompanying PHP-tool accepting OAI-PMH requests as input, processing the XML document according to the request, and returning a valid OAI-PMH response. Usage of this approach is recommended for data providers exposing metadata collections, the size of which does not justify the usage of a database.

Proposed but rejected changes. In a determined attempt to keep Version 2 of the protocol conceptually and functionally as close to Version 1.x as possible, many proposals for changes made by OAI-tech members were rejected after in-depth deliberations. Two are listed here.

Machine Readable Rights. Beginning with the first release of the OAI-PMH, there has been considerable discussion on the OAI-general list about intellectual property rights for the metadata harvested through the OAI-PMH. Because a fundamental concept in the OAI-PMH is the physical exchange of metadata, this seemed an issue that could not be ignored. Thus, the addition of machine-readable statements identifying the boundaries of acceptable use of harvested metadata was considered during the Version 2.0 process.

The decision was made not to deal with this matter at the protocol level, but rather to motivate communities to investigate it in detail and to propose solutions that could be made part of Implementation Guidelines. It was determined that emerging DRM vocabularies have not reached a required level of maturity and that acceptance of solutions in this area would very much depend on individual communities. By means of the extensible container structures, the protocol does provide locations where such information can eventually be stored.

Search functionality. A proposal to extend the functionality of the protocol by allowing simple boolean queries using Dublin Core elements as arguments was rejected. The general perception was that such functionality clearly crossed the borders of a harvesting protocol, as well as the boundaries with other standardization efforts such as ZING SRW and SRU [9].

4 Prospects for the Future

The release of v.2.0 of the OAI-PMH in June 2002, marks the end of a 17 month experimentation period that was announced with the release of v.1.0 of the protocol in January 2001. Version 2.0 of the OAI-PMH is released as a protocol ready for prime time. In order to motivate migration of existing repositories to Version 2.0, the OAI registration service will no longer accept registration of v.1.x repositories starting

September 1st 2002; repositories not conformant to v.2.0 will be removed from the registry on December 1st 2002.

Interest has been expressed in the creation of a SOAP [13] version of the OAI-PMH. This path will most likely be explored in the upcoming months. The current thinking with this respect is that the SOAP version should provide the same functionality as v.2.0 of the OAI-PMH, making it a parallel rather than a new version. As a matter of fact, care has already been taken to prepare protocol responses in OAI-PMH v.2.0 for usage in a SOAP version.

It is expected that the OAI-PMH will become part of the basic infrastructure for work in digital libraries. Early indications of this trend have become visible during the past months, as communities have moved from talking about the protocol to talking about projects in which the protocol is used, to talking about projects and failing to mention usage of the protocol. As a result, the emergence of community-specific implementations of the protocol is anticipated, with innovations expected especially in the realm of XML Schema for metadata formats, for collection and set-level description, shared set-structures and machine-readable digital rights.

A fundamental question faced by the authors, in their roles of OAI Executives, is concerned with the future of the OAI, its protocol and its registration service. The generous funding in support of OAI activities provided by the Digital Library Federation and the Coalition of Networked Information will terminate by the end of 2002. It is not in the authors' nature to run an organization for the its own sake. However, both authors share a sense of responsibility for what they have helped create, and as such they see it as their task to find an alternative strategy that provides adequate guarantees regarding the successful maintenance of the OAI-PMH and its evolution. At the time of writing of this paper, the details on how to do this are being developed.

Acknowledgements

Support for Open Archives Initiative activities comes from the Digital Library Federation, the Coalition for Networked Information, and from National Science Foundation Grant No. IIS-9817416. The authors acknowledge the fantastic support this work has gotten from the OAI Steering Committee, OAI Technical Committee, and the OAI community as a whole.

References

- [1] OAI-implementers list and archive <http://www.openarchives.org/mailman/listinfo/OAI-implementers>.
- [2] *Australian Museums & Galleries Online*, 2002 <http://amol.org.au/>.
- [3] *CIMI Consortium*, 2001 <http://www.cimi.org>.
- [4] *Dublin Core Metadata Initiative*, <http://dublincore.org>.
- [5] *eprints.org*, <http://www.eprints.org>.

- [6] *OAister...find the pearls*, 2002 <http://oaister.umdl.umich.edu/>.
- [7] *Open Archives Forum*, 2002 <http://www.oaforum.org/>.
- [8] *Open Archives Protocol for Metadata Harvesting v.2.0 scheduled for release*, 2002 <http://www.openarchives.org/news/OAI-2.0A-PR.html>.
- [9] *ZING: Z39.50 International: Next Generation*, 2002 <http://www.loc.gov/z3950/agency/zing/>.
- [10] W. Y. Arms, D. Hillmann, C. Lagoze, D. Krafft, R. Marisa, J. Saylor, C. Terrizzi, and H. Van de Sompel, "A Spectrum of Interoperability: The Site for Science Prototype for the NSDL," *D-Lib Magazine*, 8 (1), 2002.
- [11] S. Bird and G. Simons, *OLAC Protocol for Metadata Harvesting*, 2001 <http://www.language-archives.org/OLAC/protocol.html>.
- [12] S. Bird, G. Simons, and C.-R. Huang, "The Open Language Archives Community and Asian Language Resources," presented at Workshop on Language Resources in Asia, 6th Natural Language Processing Pacific Rim Symposium (NLPRS), Tokyo, 2001.
- [13] D. Box, D. Ehnebuske, G. Kakivaya, A. Layman, N. Mendelsohn, H. F. Nielson, S. Thatte, and D. Winer, "Simple Object Access Protocol (SOAP) 1.1," World Wide Web Consortium, W3C Note May 08 2000. <http://www.w3.org/TR/SOAP/>.
- [14] S. Bradner, "Key words for use in RFCs to Indicate Requirement Levels," IETF RFC 2119, March 1997. <http://www.ietf.org/rfc/rfc2119.txt>.
- [15] T. W. Cole, J. Kaczmarek, P. F. Marty, C. J. Prom, B. Sandore, and S. Shreeves, "Now That We've Found the 'Hidden Web,' What Can we Do with It? The Illinois Open Archives Initiative Metadata Harvesting Experience," presented at Museums and the Web, Boston, 2002.
- [16] N. Dushay, "Localizing Experience of Digital Content via Structural Metadata," presented at submission to Joint Conference on Digital Libraries, Portland, Oregon, 2002.
- [17] D. C. Fallside (ed.), "XML Schema Part 0: Primer," World Wide Web Consortium, W3C Candidate Recommendation CR-xmlschema-0-20001024, 24 October 2000. <http://www.w3.org/TR/xmlschema-0/>.
- [18] S. Harnad, "Free at Last: The Future of Peer-Reviewed Journals," *D-Lib Magazine*, 5 (12), 1999.
- [19] A. Kenney, N. McGovern, P. Boticelli, R. Entlich, C. Lagoze, and S. Payette, "Preservation Risk Management for Web Resources," *D-Lib Magazine*, 8 (1), 2002.
- [20] C. Lagoze, W. Arms, S. Gan, D. Hillmann, C. Ingram, D. Krafft, R. Marisa, J. Phipps, J. Saylor, C. Terrizzi, W. Hoehn, D. Millman, J. Allan, S. Guzman-Lara, and T. Kalt, "Core Services in the Architecture of the National Digital Library for Science Education (NSDL)," Cornell University, Ithaca, arXiv Report cs.DL/0201025, January 29 2002. <http://arxiv.org/abs/cs.DL/0201025>.
- [21] C. Lagoze and H. Van de Sompel, "The Open Archives Initiative: Building a low-barrier interoperability framework," presented at Joint Conference on Digital Libraries, Roanoke, VA, 2001.
- [22] S. Lawrence and C. G. Giles, "Search Engines Fall Short," *Science*, 285 (5426), pp. 295, 1999.

- [23] X. Liu, "DP9 Service Provider for Web Crawlers," *D-Lib Magazine*, 7 (12), 2001.
- [24] J. Perkins, "A New Way of Making Cultural Information Resources Visible on the Web: Museums and the Open Archives Initiative," presented at Museums and the Web, Seattle, 2001.
- [25] C. Shapiro and H. R. Varian, *Information rules : a strategic guide to the network economy*. Boston, Mass.: Harvard Business School Press, 1999.
- [26] H. Suleman, "Enforcing Interoperability with the Open Archives Initiative Repository Explorer," presented at Joint Conference on Digital Libraries, Roanoke, 2001.
- [27] H. Van de Sompel, T. Krichel, and M. L. Nelson, "The UPS Prototype: an experimental end-user service across e-print archives," *D-Lib Magazine*, 6 (2), February, 2000.
- [28] H. Van de Sompel and C. Lagoze, "The Santa Fe Convention of the Open Archives Initiative," *D-Lib Magazine*, 6 (2), 2000.
- [29] D. Waters, "The Metadata Harvesting Initiative of the Mellon Foundation," ARL, ARL Bimonthly Report 217, August 2001.
<http://www.arl.org/newsltr/217/waters.html>.
- [30] M. Wolf and C. Wicksteed, "Date and Time Forates," World Wide Web Consortium, W3C Note 15 September 1997. <http://www.w3.org/TR/NOTE-datetime>.
- [31] L. L. Zia, "Growing a National Learning Environments and Resources Network for Science, Mathematics, Engineering and Technology Education," *D-Lib Magazine*, 7 (3), 2001.