Better interoperability through the Open Archives Initiative

Michael L. Nelson

NASA Langley Research Center, Hampton Virginia 23681 USA

email: M.L.Nelson@larc.nasa.gov

The Open Archives Initiative (OAI) is an evolving protocol and philosophy regarding interoperability for digital libraries (DLs). Previously, "distributed searching" models were popular for DL interoperability. However, experience has shown distributed searching systems across large numbers of DLs to be difficult to maintain in an Internet environment. The OAI is a move away from distributed searching, focusing on the arguably simpler model of "metadata harvesting". Perhaps the strongest and distinguishing feature of OAI is its simplicity: by being "smaller" than previous interoperability projects, it actually allows for more powerful and adaptable configurations and deployments. Key concepts in OAI include the separation of responsibilities of "service provider" and "data provider" and the use of community-specific metadata sets (with Dublin Core as a lingua franca). This paper gives a brief history of the OAI, an examination of the protocol itself, and lists some of the current projects and future directions.

1. INTRODUCTION

The Open Archives Initiative (OAI) (1) is an international consortium focused on furthering the interoperability of digital libraries (DLs) through the use of "metadata harvesting". Many previous DL interoperability projects focused on "distributed searching" as the method for federating different DLs into a single service. While feasible for small numbers of nodes (e.g., < 20), large-scale distributed searching has proven difficult in an Internet environment for large numbers of nodes (e.g., > 100).

The OAI retreats from the model of distributed searching, and attempts far less technical specification than previous DL interoperability projects. As a result of this decreased scope, the OAI is proving to be a more flexible and resilient for interoperability – a sort of "RISC" (reduced instruction set computer) model for DL interoperability. The OAI defines only a generic bulk metadata transport protocol, and leaves other features to be borrowed from other technologies or implemented as independent services.

Key to understanding the philosophy of the OAI is understanding the separation of responsibilities of "service provider" (SP) and "data provider" (DP). In practice, a SP and a DP can reside in the same entity, but it is important to understand the distinction. A DP is a repository (or "archive" – the "archive" in OAI is a remnant of its e-print origins, it does not carry the technical connotation of preservation commitment that archivists reserve for the word) is simply a collection of metadata records (which may or may not point to corresponding full-text documents). A SP provides value added services (e.g., searching, browsing) on the metadata extracted from one or more DPs. Figure 1 illustrates a simple SP/DP model, where users can choose between two SPs, one of which harvests from two DPs and another that harvests from three DPs. The SPs are free to define their own services, presentation and interfaces tailored to the user base. These services could be complimentary or competitive.

The OAI metadata harvesting protocol (MHP) defines the interaction between SPs and DPs. A DL can be both an SP and DP (in fact, historically this is often true); however the separation of these functions allows some organizations to focus only on being "publishers" (i.e., populating their DP) and some organizations to focus on the development of value added services targeted to their specific customer base. The OAI MHP is quite simple, and its utility derives just as much from what is missing as what is explicitly defined. Although the OAI origins come from academic e-print distribution projects, it is applicable to a variety of applications, including the structured exposure of the "hidden web" (2) or "deep web" (3).

Figure 1. Relationship of Service Providers and Data Providers

2. BACKGROUND

The OAI grew out of the meeting surrounding the presentation of the Universal Preprint Service (UPS) DL. The UPS prototype was a feasibility study for the creation of cross-archive, end-user services that was initiated by Herbert Van de Sompel, Richard Luce and Paul Ginsparg, all then at Los Alamos National Laboratory (LANL). With the premise that users would prefer access to a federation of DLs instead of the disparate islands of DLs that existed, they proposed a project aimed at identifying the key issues in actually *creating* an experimental end-user service for data originating from existing production archives. This included almost 200,000 records harvested from six popular production DLs. The implementation of the UPS was coordinated by Herbert Van de Sompel, Thomas Krichel and Michael Nelson, and a full discussion of the history, creation and demonstration of UPS project can be found in (4).

Some of the constraints of the UPS prototype influenced what would become the OAI. First, UPS was developed under a tight deadline (four months), which naturally narrowed the scope of the project. Secondly, neither resources (storage space, time) nor permission was available to download the full text of the six participating DLs. This resulted in UPS focusing only on the metadata of the records, not the full-text records themselves. Thirdly, the coordinators had control of only two of the DLs

used in UPS – there could be no significant demands placed on the DLs; UPS itself would have to address any shortcomings that the DLs might have. This influenced the OAI philosophy that it should be as simple as possible to be a DP and SPs should bare the burden of any complexity. Lastly, the UPS was also used as a demonstration for the SFX reference linking service (5) and "bucket" digital object technology (6). This emphasized the focus of providing services on metadata harvested from other archives.

During the initial meeting in Santa Fe, New Mexico in October 1999, the Santa Fe Convention emerged (7). Originally, this was a subset of the Dienst DL protocol (8), and a separate minimal metadata format was defined for baseline interoperability. In September 2000, a meeting in Ithaca New York defined the Santa Fe Convention, later renamed the OAI metadata harvesting protocol, and the Dienst subset and OAI specific metadata set were exchanged for an OAI specific protocol and unqualified Dublin Core (DC) (9), respectively. An alpha test group of DPs and SPs were formed in November 2000, and the OAI MHP was released at public meetings in January 2001 in the United States, and February 2001 in Europe (1).

3. DIGITAL LIBRARY INTEROPERABILITY MODELS

Interoperability is a difficult problem, and previous attempts to address it have resulted in an escalating level of complexity. Eventually, the complexity of interoperability solutions began to rival the problem itself. The OAI can be considered a huge step backwards in scope, ambition and focus in comparison with previous projects. Earlier attempts at repository-level interoperability had been made, but had generated "limited consensus" regarding specific functionality (10). Perhaps such attempts were simply premature within the community, and the OAI simply occurred at the right time.

An early champion of distributed searching in web DLs was the Dienst protocol, which was used to construct the Networked Computer Science Technical Reference Library (NCSTRL) (8). Dienst defines a rich suite of services for user interfaces, repositories, searching, and collections. It also has a built-in document object model and, for example, handles scanned documents nicely. At its peak, NCSTRL had over 100 university computer science departments (along with a few governmental and industrial research laboratories), each maintaining their own Dienst server. The initial architecture involved broadcasting each user query to all nodes. Over time, the fully distributed search model in Dienst became a distributed / centralized search hybrid, with regional meta-servers, master meta-servers and connectivity regions introduced to retain acceptable response to user queries (8). Powell & French also studied the growth and reliability of NCSTRL and found the "[r]eliability of the distributed system is low" (11).

Another distributed searching interoperability project was the STARTS protocol (12). STARTS defined and introduced solutions for the three problems that are common to distributed searching:

1. the "source–metadata" problem
(where should the queries be sent?)
2. the "query–language" problem
(what query syntax and semantics does a DL expect?)
3. the "rank–merging" problem
(how to meaningfully merge the result sets from more than one DL?)

STARTS defined methods for addressing each of these problems, by allowing a search service to export generalized descriptions of itself (the service's nature and purpose), its technical capabilities (what syntactic features it supports), and its contents (the frequency and distribution of terms it has indexed). Since STARTS was aimed at the commercial web search companies, it had the additional capability of providing its self description without compromising the proprietary nature of its collection and ranking algorithms. Though quite powerful, STARTS was never widely implemented, presumably due in part to its complexity.

Partially similar to STARTS, another approach to distributed searching involves third parties constructing descriptions of the capabilities of (potentially uncooperative) DLs, and providing searching services based on these descriptions. Two examples include the Lyceum project (13) and the InterOp project (14). The former provided a simple web interface for describing the capabilities of DLs and storing these descriptions at a central site, and the latter provided a more sophisticated extensible markup language (XML) encoded implementation. While both projects provide passable solutions to the source–metadata and query–language problems, they are insufficient to address the rank–merging problem.

No discussion of interoperability would be complete without mention of Z39.50 (15). While present in commercial information systems, Z39.50 appears to have made little actual impact in the web DL community. Presumably this is because the complexity of Z39.50 is significant, and many DLs come from simple, grass–roots origins (16). Regardless of reason, it is worth noting that the OAI protocol was crafted with assistance of those extremely knowledgeable about Z39.50: Clifford Lynch (Santa Fe, Ithaca) and Ray Dennenberg (Ithaca). Although the protocols can be used to create similar services (e.g., federated searching services), Lynch has stated that the two protocols are "meant for different purposes, with very different design parameters", "neither is a substitute for the other" and the DL community would do well to think about "bridges and gateways" (17).

The notion of metadata harvesting within the web context was first introduced in the Harvest system, developed at the University of Colorado (18). Harvest was actually a complex suite of tools that included gatherers, brokers and other services rolled into a single package. Although it generated significant interest, it was quickly absorbed into a variety of commercial products, and the development of the open source version suffered, limiting its deployment and adoption.

4. THE OAI METADATA HARVESTING PROTOCOL

The OAI metadata harvesting protocol defines six "verbs" that comprise the communication between data providers and service providers (Table 1). Three of verbs, Identify, ListSets, and ListMetadataFormats provide metadata about the archive itself. They provide a description of the archive, points of contact (POC),

pointers to any usage policy that might be in effect, and lists some of the OAI conventions that may be in use in the archive (sets, metadata formats).

The real work in metadata harvesting is accomplished by the verbs ListIdentifiers, GetRecord and ListRecords. The identifiers in a repository refer to the metadata records themselves – not any full-text content that may exist. Identifiers are hierarchically constructed, beginning with "oai:", followed by the repository identifier, followed by a locally defined string. ListIdentifiers and ListRecords also take arguments that allow specification of dates and date ranges, allowing for incremental harvesting (e.g., "harvest all records added since 2001-05-31"). All three verbs take optional set arguments (see below) (e.g. "harvest all records with "subject=physics and status=approved"). GetRecord and ListRecords also take arguments to request the desired metadata format (note that the DP may or may not have the requested format to disseminate). All of the arguments listed are orthogonal, and can be combined to narrow the requested range of documents.

<i>Verb</i>	<i>Description</i>
Identify	returns a description of the archive (name, POC, etc.)
ListSets	returns a list of sets in use by the archive
ListMetadataFormats	returns a list of metadata formats used by the archive
ListIdentifiers	returns a list of record ids (possibly matching some criteria)
GetRecord	given a record id, returns that record
ListRecords	returns a list of records (possibly matching some criteria)

Table 1. OAI Metadata Harvesting Protocol Verbs

All OAI verbs are passed from a SP (frequently called a "harvester" in this context) to a DP by encoding the request in an http URL, and the response from a DP is encoded in XML. If a the "base URL" of a DP is "http://naca.larc.nasa.gov/oai/"

`http://naca.larc.nasa.gov/oai/?verb=ListIdentifiers&from=2001-10-31`

then the response would be:

```
<?xml version="1.0" encoding="UTF-8"?>
<ListIdentifiers
  xmlns="http://www.openarchives.org/OAI/1.1/OAI_ListIdentifiers"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://www.openarchives.org/OAI/1.1/OAI_ListIdentifiers
    http://www.openarchives.org/OAI/1.1/OAI_ListIdentifiers.xsd">
  <responseDate>2001-11-14T23:23:57+00:00</responseDate>

  <requestURL>http%3A%2F%2Fnaca.larc.nasa.gov%2Foai%2Findex.cgi%3Fverb%3DListIdentifiers
%26from%3D2001-10-31</requestURL>

  <identifier>oai:NACA:1958:naca-rm-a58c03</identifier>
  <identifier>oai:NACA:1958:naca-rm-a58c21</identifier>
  <identifier>oai:NACA:1958:naca-tm-1441</identifier>
</ListIdentifiers>
```

which lists three identifiers that have been added (or changed) since October 31, 2001. A harvester could then either issue three separate GetRecord verbs to retrieve the new records, or issue a single ListRecords verb with the same "from" argument to retrieve the full metadata records. It is important to note that all of the OAI transactions are not meant to be directly observed by the end-user. OAI is middleware and should be completely transparent to the end-user. For developers that are adding an OAI interface to their DL, the OAI website has a number of tools and code bits to assist in the task. Of particular importance is the Repository

Explorer (19), which allows users to interactively test the performance and conformance of their DP.

There are a number of optional concepts within the OAI MHP:

- sets: the use of sets is optional, and if implemented, it is left entirely up to the DP to define/adopt a suitable set structure
- parallel metadata: unqualified Dublin Core is the "default" metadata format and a DC dissemination must always be available; however the adoption of richer, community-specific metadata formats is encouraged
- flow control: it is possible for a DP to control the rate of requests that it receives from a harvester through the use of an optional, locally defined "resumptionToken"

For a more lengthy discussion of some of these concepts, see Warner (20), Lagoze & Van de Sompel (21) or the OAI Protocol document itself (22). The goal behind making these concepts optional was to keep the barrier for becoming a DP as low as possible. However, DPs that wish to engage in more sophisticated negotiation with harvesters have the semantics to express their capabilities and desires.

5. OBSERVATIONS

There are a number of misconceptions about the OAI, and these frequently have to do with attributing too much to the scope of OAI. First, it is important to remember that the OAI is not a DL by itself, it is always an interoperability layer, or "front-end", that is added to existing DLs. As such, the protocol has no mechanism for inputting and deleting items – this is handled by the native DL, not by the OAI interface.

Furthermore, the OAI MHP protocol has no mechanisms to support distributed searching. While support for date ranges is viewed as necessary for supporting hierarchical harvesting and the optional set construct is present, the OAI MHP does *not* support keyword searching, argument wildcards, filters, or other advanced, repository-side search capabilities. Such facilities would violate the OAI philosophy of making the DP as simple as possible, with the SP responsible for picking up the slack.

Another realm that is outside the scope of the protocol is terms and conditions. The OAI MHP has no authentication mechanisms, no access control lists or anything similar. Rather than build this into the protocol, it relies on the mechanisms of the transport protocol, http. OAI interfaces have been built for DLs featuring varying levels of restriction, only allowing harvesting from certain Internet addresses or require usernames / passwords. The "Open" in Open Archives Initiative refers to the openness of the protocol and allowing a repository to expose its contents in a well-defined manner. "Open" does not carry any requirements regarding the intellectual property restrictions of the metadata.

Similarly, concepts such as OAI mirrors, caches, and other OAI interfaces (peer, master/slave) are not defined in the protocol. These capabilities are possible, of course, but are again deferred to the transport protocol or to separate higher-level

services. Indeed, some projects have already demonstrated a number of services that lie outside the protocol, such as load balancing through redirection, hierarchical harvesting, and caching of full-text content. And while much has been made regarding the difference between metadata harvesting and distributed searching models, it is easy to imagine hybrid architectures, perhaps providing distributed searching across a small number of homogeneous nodes (e.g., ten) that have hierarchically harvested their metadata from many nodes (e.g. thousands – or more).

6. CURRENT PROJECTS

Any list that aimed for completeness while enumerating OAI developments and projects would soon be out of date. Readers are encouraged to check the OAI website for complete and breaking news. Cornell University runs a central registry for DPs and SPs that is not required for OAI operation, but is frequently used to advertise existence. Currently, three service providers and 41 data providers are registered (twelve in Europe, one in Mexico, and the rest in the U.S.). However, at least four more SPs are known to be in development or testing at this point, with more expected in the near future. It is also known that many other DPs are in development, and some are in "private" operation (they are being used internally at their respective organizations and will never be centrally registered). One of the DPs is a commercial publisher, with several other publishers known to be experimenting with OAI interfaces.

For the aspiring DP, a number of tools and code bits are available from the OAI website for grafting an OAI interface to a number of popular databases, directory servers and filesystems. For those who are installing a DL from scratch, a number of options exist. Eprints.org (23) offers a complete, stand-alone departmental size DL system with an OAI interface. For those desiring a smaller foot-print DP, Kepler (24) is a "personal-size" DP, suitable for home PCs or laptops. Because Kepler is expected to be deployed on systems where connectivity is not always assured, a full-text content caching service is implemented in addition to the OAI MHP.

As for SPs, in addition to the Repository Explorer already mentioned, there is Arc (15), which provides a demonstration federation service over all known DPs and currently indexes over 1 million metadata records. In addition, Arc also features hierarchical harvesting: even though it is a SP, it can also function as a DP, allowing other harvesters to extract the local copies of records that Arc has harvested from DPs. Torii (26) is a SP that provides an alternate interface to the records in arXiv (27) as well as the local DP of the International School for Advanced Studies of Trieste (28).

Some SPs currently under development include a new, OAI-based implementation of NCSTRL (29). The University of Southampton is working on cite-base (30), a SP that extracts citations from full text of articles described in OAI records and provides searching on the result. The Open Language Archives Community (31) is also working on a SP specific to the linguistics community.

7. FUTURE DEVELOPMENTS

Believing that a stable protocol would encourage adoption, the the OAI Steering Committee has frozen the OAI MHP since its debut in January 2001. Drawing from the OAI community's experience to date, the OAI Technical Committee is currently

revising the protocol, with an expected completion date of May 2002. The protocol will incrementally evolve, with the addition of a handful of features and clarification of some existing concepts. However, the protocol will remain fundamentally unchanged, as well as the OAI commitment to low-barrier interoperability, with the bias toward easy creation of DPs.

Furthermore, the community is still gaining experience with metadata harvesting for DL interoperability. Metadata harvesting may prove simpler than distributed searching, but there still a number of subtle, yet difficult issues regarding metadata harvesting. Data is still being collected on issues such as optimal frequency of harvests, data coherency in distributed environments, how to handle deletions, and tuning of flow control mechanisms. While much work has been done in the general web community on some of these issues, it remains to be seen if some mechanism, such as the freshness metrics of Cho & Garcia-Molina (32), are applicable in a DL environment.

Finally, it is hoped that communities will continue to adopt richer metadata sets. Dublin Core is attractive in that it provides a lingua franca appropriate for resource discovery, but sophisticated SPs will need richer metadata to provide better services. This can be achieved through qualified DC, or through separate metadata formats altogether. As long as an XML schema is available that define the format, it can be used in the OAI MHP. In addition to DC, other formats currently in use include MARC (33), RFC-1807 (34), Open Languages Archives Community Metadata Set (35), and the Electronic Theses and Dissertation Metadata Set (36).

8. CONCLUSIONS

From the point of view of the end-user, the perfect OAI implementation is advertised, not seen. The OAI metadata harvesting protocol is a generic bulk metadata transport that has generated significant international interest as a tool for DL interoperability. The strength of the protocol is as much what it leaves out as what it defines. It utilizes other technologies when possible (http, XML schemas, Dublin Core), and defines its own features when necessary. The OAI embraces a philosophy that favors the easy creation of data providers, and requires the service providers to "work harder" to compensate for irregularities or missing functionality in data providers. The OAI MHP focuses only on metadata, not full-text, and is always a front-end to an existing DL. While experience is still to be gained working with the metadata harvesting paradigm, it is expected that it will yield greater flexibility and interoperability than the previously more popular model of distributed searching.

REFERENCES

1. <http://www.openarchives.org/>
2. RAGHAVAN, S. and GARCIA-MOLINA, H. Crawling the hidden web. *Proceedings of the 27th VLDB conference*, 2001, 129-138.
<http://www-db.stanford.edu/~rsam/pubs/vldb01/vldb01.pdf>
3. BERGMAN, M. K. The deep web: surfacing hidden value. *Journal of Electronic Publishing*, 7(1), 2001.
<http://www.press.umich.edu/jep/07-01/buckholtz.html>

4. VAN DE SOMPEL, H., KRICHEL, T., NELSON, M. L., HOCHSTENBACH, P., LYAPUNOV, V. M., MALY, K., ZUBAIR, M., KHOLIEF, M., LIU, X. and O'CONNEL, H. The UPS prototype: an experimental end-user service across e-print archives. *D-Lib Magazine*, 6(2), 2000.
<http://www.dlib.org/dlib/february00/vandesompel-ups/02vandesompel-ups.html>
5. VAN DE SOMPEL, H and HOCHSTENBACH, P. Reference linking in a hybrid library environment. Part 2: SFX, a generic linking solution. *D-Lib Magazine*, 5(4), 1999.
http://www.dlib.org/dlib/april99/van_de_sompel/04van_de_sompel-pt2.html
6. NELSON, M. and MALY, K. Buckets: smart objects for digital libraries. *Communications of the ACM*, 44(5), 2001, 60–62.
7. VAN DE SOMPEL, H and LAGOZE, C. The Santa Fe Convention of the Open Archives Initiative. *D-Lib Magazine*, 6(2), 2000.
<http://www.dlib.org/dlib/february00/vandesompel-oai/02vandesompel-oai.html>
8. DAVIS, J. R. and LAGOZE, C. NCSTRL: design and deployment of a globally distributed digital library. *Journal of the American Society for Information Science*, 51(3), 2000.
9. WEIBEL, S. L. and LAGOZE, C. An element set to support resource discovery. *International Journal on Digital Libraries*, 1(2), 1997, 176–186.
10. SCHERLIS, W. L. Repository interoperability workshop: towards a repository reference model. *D-Lib Magazine*, 2(10), 1996.
<http://www.dlib.org/dlib/october96/workshop/10scherlis.html>
11. POWELL, A. L. and French, J. C. Growth and server availability for the NCSTRL digital library. *Proceedings of the 5th ACM Conference on Digital Libraries*, 2000, 264–265.
12. GRAVANO, L. CHANG, C.-C. K., GARCIA-MOLINA, H. and PAEPCKE, A. STARTS. *Proceedings of the 1997 ACM SIGMOD international conference on management of data*, 1997, 207–218.
13. MAA, M-H., ESLER, S. L., and NELSON, M. L. Lyceum: a multi-protocol digital library gateway. *NASA Technical Memorandum 112871*, 1997.
<http://techreports.larc.nasa.gov/ltrs/PDF/1997/tm/NASA-97-tm112871.pdf>
14. ZUBAIR, M., MALY, K., AMEERALLY, I. and NELSON, M. L. Dynamic construction of federated digital libraries. *Poster, WWW9 conference*, 2000.
<http://www.www9.org/final-posters/poster17.html>
15. <http://www.loc.gov/z3950/agency/>
16. ESLER, S. L. and NELSON, M. L. Evolution of scientific and technical information. *Journal of the American Society for Information Science*, 49(1), 1998, 82–91.

17. LYNCH, C. A. Metadata harvesting and the Open Archives Initiative. *ARL Bimonthly Report* 217, 2001.
<http://www.arl.org/newsltr/217/mhp.html>
18. BOWMAN, C. M., DANZIG, P. B., HARDY, D. R., MANBER, U., SCHWARTZ, M. F. and WESSELS, D. P. Harvest: a scalable, customizable discovery and access system. *University of Colorado Computer Science Technical Report CU-CS-732-94*, 1994.
<ftp://ftp.cs.colorado.edu/pub/cs/techreports/schwartz/Harvest.Jour.ps.Z>
19. SULEMAN, H. Enforcing interoperability with the Open Archives Initiative Repository Explorer. *Proceedings of the 1st ACM/IEEE Joint Conference on Digital Libraries*, 2001, 63–64.
20. WARNER, S. Exposing and harvesting metadata using the OAI metadata harvesting protocol: a tutorial. *HEP Libraries Webzine*, 4, 2001.
<http://library.cern.ch/HEPLW/4/papers/3>
21. LAGOZE, C. and VAN DE SOMPEL, H. The Open Archives Initiative: building a low-barrier interoperability framework. *Proceedings of the 1st ACM/IEEE Joint Conference on Digital Libraries*, 2001, 54–62.
22. http://www.openarchives.org/OAI_protocol/openarchivesprotocol.html
23. <http://www.eprints.org/>
24. MALY, K., ZUBAIR, M. and LIU, X. Kepler: an OAI data/service provider for the individual. *D-Lib Magazine*, 7(4), 2001.
<http://www.dlib.org/dlib/april01/maly/04maly.html>
25. LIU, X., MALY, K., ZUBAIR, M. and NELSON, M. L. Arc: an OAI service provider for digital library federation. *D-Lib Magazine*, 7(4), 2001.
<http://www.dlib.org/dlib/april01/liu/04liu.html>
26. <http://torii.sissa.it/>
27. <http://www.arxiv.org/>
28. BERTOCCO, S. Torii, an open portal over open archives. *HEP Libraries Webzine*, 4, 2001.
<http://library.cern.ch/HEPLW/4/papers/4/>
29. <http://www.ncstrl.org/>
30. <http://cite-base.ecs.soton.ac.uk/>
31. <http://www.language-archives.org/>
32. CHO, J. and GARCIA-MOLINA. Synchronizing a database to improve freshness. *Proceedings of the 2000 ACM International Conference on Management of Data*, 2000, 117–128.
<http://www-db.stanford.edu/~cho/papers/cho-synch.pdf>

33. <http://www.loc.gov/marc/>
34. LASHER, R. *and* COHEN, D. A format for bibliographic records. *Internet RFC-1807*, 1995.
<ftp://ftp.isi.edu/in-notes/rfc1807.txt>
35. <http://www.language-archives.org/OLAC/olacms.html>
36. <http://www.ndltd.org/standards/metadata/>