

The Current State of Linked Data in Libraries, Archives, and Museums

Since the last issue of *Library Technology Reports* (*LTR*) on Linked Data (LD) in July 2013, the library, archive, and museum (LAM) communities have put considerable work into developing new LD tools, standards and published vocabularies, as well as explored new use cases and applications. In 2013, there was already a range of LD systems in production, and in the past two years, the number of systems has grown steadily. Alongside this growth and experimentation, the discussion of Linked Data and Linked Open Data (LOD) has explored the nuanced differences between schemas such as BIBFRAME and BIBFRAME Lite, has explored the expansion of vocabularies and technologies, and has expanded around themes of technology adoption, LD literacy, evolution of standards and schemas, case studies in adoption, and studies of value and impact.

The 2013 *LTR* issue on LD used a largely technical lens to explore these issues, as there were many unanswered questions about how LAM organizations might apply emerging LD concepts in their metadata and information systems. In studying three important LD platforms (Europeana, OAI-PMH, and DPLA) and in devoting a chapter to exploring the fundamentals of LD, that issue sought to capture the state of adoption and technology use across the LAM community. This update on LD adoption takes a different approach by exploring at a broader level the issues, trends, and LD programs that are shaping our community perspectives. In order to do this, chapter 1 of this issue considers the broad state of LD adoption. Chapter 2 examines projects, services, and research efforts with a goal of better understanding the overall trajectory of adoption. Chapter 3 takes a more detailed look at the vocabularies, schemas, standards, and technologies that are forming the foundation of LD, and chapter 4

considers the policies and practices that are influencing the community and considers next steps that may hold promise in the LAM community.

In order to paint a picture of current efforts and adoption in Linked Data as well as to project the potential future of LD efforts, this issue draws on surveys of LD adoption, updates from national and international project teams, and selective exploration of technical topics that are emerging as new concepts in LD and are likely to influence LD adoption in the coming year. Just as with the 2013 issue, this update serves two purposes. First, it seeks to collect project reports and literature to synthesize ideas and trends as well as inform perspectives on the current state of LD adoption. Second, this issue seeks to capture and document current thinking and practice in LD, recognizing that, at this point, LD has become part of the central discourse in LAM communities, influencing the education and operating principles of the information professions.

The State of Linked Data Adoption

This section examines the findings of a 2014 survey on LD adoption, considers technical developments around LD in LAM contexts specifically, considers how projects and standards are evolving, and discusses broadly the visibility and maturity of projects.

Survey Results from LD Adoption

In 2014, OCLC staff conducted a survey on LD adoption, a survey that is being repeated for 2015. The analyzed results from the 2014 survey are captured in a series of blog posts on the site hangingtogether.org

and provide a substantial window into the state of LD deployment in LAM institutions.¹ The survey surfaced 172 projects, of which 76 included substantial description. Of those 76 projects, over a third (27) were in development. The larger, in terms of metadata transformed, projects included OCLC's WorldCat.org, Library of Congress's (LoC) id.loc.gov service, and the British Library's British National Bibliography.² General descriptions of selected projects are available in the second blog post as well as the raw data from the survey.³ A revised survey closed in August 2015 and results, although not available at the time of this writing, should be available on the OCLC Linked Data Research web page by the date of publication.

OCLC Linked Data Research

www.oclc.org/research/themes/data-science/linkedata.html

One interesting area of analysis from the 2014 survey focused on intended use cases and overall purpose of a LD project. Common use cases cited included “enrich[ing] bibliographic metadata or descriptions,” “interlinking,” “as a reference source and to . . . harmonize data from multiple sources,” “[to] automate authority control,” “[to] enrich an application.”⁴ In addition, the most common reasons for creating an LD service were to publish data more widely and to demonstrate potential use cases and impact.⁵ In addition, the Linked Data for Libraries (LD4L) group has gathered a set of use cases to inform their work.⁶ These use cases have been clustered into six main areas including “Bibliographic + Curation” data, “Bibliographic + Person” data, “Leveraging external data including authorities,” “Leveraging the deeper graph,” “Leveraging usage data,” and “Three-site services” (e.g., enabling a user to combine data from multiple sources).

Although the analyzed data from the survey showed that a wide range of vocabularies were used in the projects reported, there was also a strong cluster around just a few published vocabularies. According to Smith-Yoshimura, the most commonly used LD data sources were id.loc.gov, DBpedia, GeoNames, and VIAF.⁷ Data in the projects analyzed was often bibliographic or descriptive in nature. As captured in the analysis by Smith-Yoshimura, the most common organizational schemas used were Simple Knowledge Organization System (SKOS), Friend of a Friend (FOAF), Dublin Core and Dublin Core terms, and Schema.org.⁸ In addition to this short list of highly used vocabularies and schemas, the data shows a much longer list of all of the vocabularies cited in the results.

The analyzed results of the survey indicated that Resource Description Framework (RDF) serialized in the eXtensible Markup Language (XML) was commonly used, as was RDF serialized in JavaScript

Object Notation (JSON) and Terse RDF Triple Language (Turtle).⁹ Advice from implementers, the content of the sixth blog post on the LD survey, presents a range of perspectives on project management, project scope, and possible technologies and standards to use in development.¹⁰ One sentiment captured in the results is the importance of publishing “useful” data. This sentiment is part of the LOD building blocks popularized by Berners-Lee, especially the rule “When someone looks up a URI, provide useful information, using the standards.”¹¹ This notion, although seemingly obvious, has become part of subsequent recommendations around the creation of LD. For example, the CIDOC Conceptual Reference Model Special Interest Group (CRM-SIG) has codified this sentiment in a series of guidelines for creating and publishing LOD.¹² Of equal importance but with less guidance is the issue of data licensing. The referenced CIDOC recommendation focuses largely on technical issues and does not mention licensing recommendations. Somewhat surprisingly, in the OCLC survey results, there was a range of approaches to licensing of data, including many Creative Commons CC0 licenses but also Open Data Commons (ODC) and noncommercial use licenses.¹³ Such variation in licensing may not be a substantial issue, but it does add a level of complexity when considering what uses an organization can make of published data.

A related policy question surfaced in this survey is how LAM institutions should approach LD production or adoption. It appears that despite the transition to Linked Data for large-scale and core services such as the transformation of library MARC platforms and the migration of EAD finding aids, the community has not yet distilled a set of activities or systems into an “easy-to-implement” platform or adoption approach. Indeed, LD efforts might still be categorized as existing in the startup phase of a technology adoption hype cycle given the variation in standards, tools, approaches, and perceived benefits documented in survey results and published literature. At the same time, however, LD services have expanded to a point where they may soon reach critical mass in enabling widespread use in the LAM community. This is demonstrated in part by the continued growth of LD adopters and test programs that are working with data that would impact a large number of libraries and archives. It is also indicated by the growth of the number of triples published by these services, showing that the automation and refinement tools needed are reaching a level of maturity and that successive LD projects have more to build on.

Activities across US Libraries

Another useful source of information about developments and projects in LD is the annual updates of

research libraries in conjunction with the American Library Association (ALA) ALCTS Technical Services Directors of Large Research Libraries Interest Group.¹⁴ The fifteen public reports from June 2015 show a range of LD efforts in these libraries. For example, many institutions are pursuing education for staff via the Library Juice Academy certificate program (<http://libraryjuiceacademy.com>) or the Zepheira LibHub early adopters training (<http://zepheria.com/solutions/library/training>). Many of the reports indicate that institutions have approached LD from an exploration and research perspective (e.g., formation of a project team; establishing broad goals; working with available tools and standards to explore impact in the local environment). Trends in these reports included exploring how to leverage LD and LD URIs in discovery systems generally and potentially in local catalog applications.

Within this research thread there are a number of specific projects. As a partner in the LD4L project, Cornell has been active in an ontology group and working to set up a Vitro instance for LD cataloging.¹⁵ The Library of Congress reported its multifaceted work in BIBFRAME, providing a window into the development and testing of this schema. The report indicates that LoC is using the MarkLogic platform for development of BIBFRAME and leveraging the vocabularies at the LoC Linked Data Service Authorities and Vocabularies web page. It is projecting a test of this platform for late summer and early fall of 2015, the goal of which is to explore the application of BIBFRAME and these vocabularies in a real-world setting.¹⁶ Likewise, the National Library of Medicine (NLM) has undertaken considerable testing and development with LD, as reported elsewhere in this issue. This work includes releasing Medical Subject Headings (MeSH) as RDF. This data is being made available as annually updated downloadable files.¹⁷ Although much of the work in LD in the LAM community comes from bibliographic roots there is evidence of a growing interest in other data sources and applications. For example, in addition to traditional resource-based metadata, some institutions are working with ORCID identifiers as a way to better capture research productivity for faculty and graduate students.

VIVO

<http://vivoweb.org>

LoC Linked Data Service: Authorities and Vocabularies

<http://id.loc.gov>

Medical Subject Headings (MeSH)

<http://id.nlm.nih.gov/mesh>

In addition to containing specific project information on LD, there are several projects that seem poised to benefit from advances in LD. Migration of libraries from either older versions of their ILS or to a new open-source ILS platform (e.g., the Open Library Environment) was mentioned in a number of these reports, either as an accomplishment in 2015 or as an upcoming project in 2016. Likewise, the deployment or enhancement of discovery platforms remained a central activity. One trend, tangentially related to LD, was the publication of digital objects with open-access licenses. The University of Pennsylvania, for example, released OPenn, a resource focused on making cultural heritage materials available under Creative Commons licenses.¹⁸ With a similar goal, the University of Michigan released the Special Collections Image Bank with the goal of capturing digitized images and making them available under the appropriate license.¹⁹ These released products suggest potential paths of new development in LD, particularly the potential of these open digital platforms to enable more extensive discovery and reuse of resources and metadata.

OPenn

<http://openn.library.upenn.edu>

University of Michigan Special Collections Image Bank

<http://quod.lib.umich.edu/s/sclib>

Linked Data Trends: Technical, Application, and Visibility

Technical Developments in LD Adoption

In the past two years, the LD community has continued to focus on RDF and has increased its use of JSON serializations of RDF. Several important standards have seen increasing adoption, including the final specification of HTML5 and the definition of the RDF 1.1 standard in 2014.²⁰ HTML5 provides enhanced support for geolocation services, application cache and local data, server sent events (i.e., automatic updates from the server to the client), and support for web worker application programming interfaces (APIs; e.g., JavaScript running in the background of the client application). These interactivity tools are enabling the development of a new generation of interaction and data-rich web services and allow the web client to make extensive use of published open data. Similarly, the RDF 1.1 standard expands the utility of RDF by adding much-needed support for RDF datasets, a collection of RDF graphs, expansion of data types,

and new definitions for handling of internationalized resource identifiers (IRIs) and literals.²¹

The RDF 1.1 primer explores these concepts in more detail, in addition to providing an overview of emerging serialization languages including TriG, N-Quads, and JSON-LD.²² Each of these serialization techniques provides expanded support for named graphs, TriG extending Turtle to add this functionality and N-Quads extending N-Triples. JSON-LD, like JSON in general, has been an emerging and popular serialization platform for several years. At the same time, the increased emphasis on JSON-LD is not without controversy in the LD community. JSON has been praised for being a lightweight, platform-integrated approach but also criticized for not supporting the complex models and relationships that can be expressed in XML.²³ At the time of this writing, JSON-LD's inclusion of new keywords (e.g., @graph) has helped provide more robust support for the representation of RDF in JSON. In addition, as any casual user of LD applications in LAM contexts will observe, JSON-LD is increasingly common, featured in a number of LD enabled services including DPLA's API. Given the increasing use of JSON and JSON-LD, it is likely that the LD community would benefit from the further support of JavaScript and server integration coming from the HTML5 community.

In addition to efforts in the LD community to transform bibliographic and other metadata services and data stores (e.g., BIBFRAME, BIBFRAME Lite, Schema.org), there is considerable work being done to leverage LD to develop new products and services. Jason Clark and Scott Young, for example, recently explored the use of JSON-LD in creating and structuring e-book content.²⁴ Their work drew on several of the perceived benefits of LD creation, including search engine optimization, connection with social media networks, and connection to other resources through links and content integration. On the theme of service integration through structured and linked metadata, Suzanna Conrad explored the use of Google Analytics to study use of DSpace metadata fields.²⁵ Finding that the tag manager tool in Google Analytics was a good fit for tracking metadata fields in DSpace, Conrad pointed to an analytical application of data linking, even if the tools discussed do not surface metadata in a conventional LD platform.

Another important area of work in LD is the application of existing tools to improve the quality of data. Although not necessarily focused on generating LD, the increase in use of these tools is important to the long-term viability of data cleanup and normalization. Donnelley, for example, used a combination of Python and OpenRefine tools to clean up and normalize zip code information.²⁶ Such a task is often one of many steps that occur prior to the publication of data and is particularly important in the generation of unique pointer information such as zip code data.

This article in particular provides useful instructions in the detailed work required for such a task.

Coming from a different perspective, Bianchini and Willer explored the role of historic library standards such as International Standard Bibliographic Description (ISBD), asking how the concepts in ISBD fit with Semantic Web needs.²⁷ Their article explored a notion that is common in other areas of research around metadata standards: that our older vocabularies and approaches are not always easily mapped onto new technologies and use cases. In particular, Bianchini and Willer explored the shifting notion of resource from ISBD to the concept of a resource in RDF. Dunsire conducted a parallel analysis of ISBD and ISBD punctuation, finding similar challenges in employing this standard in semantic contexts without some level of modification.²⁸ These two works focusing on standard alignment with an emphasis on the role of older standards in new LD settings are representative of larger discussions in the LD community. The ALA Metadata Standards group, for example, has also debated the perceived value of ISBD in LD settings and recently drafted a series of guidelines for assessing metadata standards to help shape this discussion at a broader level.²⁹

Although much of the LD focus of the LAM community is on transformation of bibliographic and collection (e.g., MARC and EAD) schemas, there is also interest in authorities and translation of LD schemas to new domains. The electronic thesis and dissertation (ETD) community, for example, has looked at some level at the influence of LD models on connecting ETD repositories and enabling new scholars to enjoy more visibility on the web.³⁰ Likewise, emerging researcher ID platforms such as ORCID, ResearcherID, arXiv, Author Claim, and Scopus Author ID are pushing more communities toward LD-related discussions through the thread of name disambiguation and author-based graphs. The emergence of scholar identifiers in LD standards focused on earlier stages in an academic's career could have considerable impact in increasing awareness around LD issues (e.g., disambiguation, persistent identifiers, open data, and metadata) in the broader research community. The extent to which the maturity of the tools and the abilities of researchers and practitioners are at a state to support widespread adoption is yet to be seen, but such advances bode well for the broad appeal of LD and other Semantic Web technologies

ORCID

<http://orcid.org>

ResearcherID

www.researcherid.com

Focused more closely on enterprise tools and projects, a growing area of research seeks to advance understanding of potential systems based on services provided by DPLA, Europeana, and WorldCat. One example of this is Péter Király's work implementing translation services for queries with the goal of enabling a user to query terms across multiple languages simultaneously.³¹ In addition to work focused on exploring adaptive ways of using LD via APIs, other efforts continue on vocabulary improvement and publishing. Toves and Hickey recently documented expanded algorithms for processing dates in VIAF, demonstrating that the new approach has led to considerable improvements in normalization in the dataset.³² In a similar thread, some libraries are branching into their own targeted vocabulary creation. Hanson documented North Carolina State University's efforts to develop an LD dataset of organization names.³³ This project, having been in production for many years, is used to manage name information in library information systems and is also part of the Global Open Knowledgebase (GOKb). Each of these vocabularies represents highly impactful projects occurring at different scales in the LAM community.

NCSU Libraries, Organization Name Linked Data
www.lib.ncsu.edu/ld/onld

Global Open Knowledgebase
<http://gokb.org>

Occurring somewhat in contrast to these efforts to generate more LD or improve LD quality, there is also a strong thread of research around the use of APIs. Perhaps ironically, APIs are usually seen as a stopgap measure that is required when LD is not available, but in many cases they are the tools that enable the creation of LD in the first place. Reese, for example, completed an in-depth introduction to tools, techniques, and output associated with the WorldCat API.³⁴ Similarly, Nugraha introduced MariaDB, a replacement open-source server similar to MySQL and Sphynx, a full-text search platform that works in concert with relational databases.³⁵ While such work is more related to rather than directly connected with LD work, advances in the tools and techniques from work like this are important to laying the groundwork and making better use of available information systems.

Evolution of Projects and Standards

In the past year, the Library of Congress and OCLC have completed a report comparing their two approaches to LD creation,³⁶ while other efforts have

spawned BIBFRAME Lite, Zepheria's extended BIBFRAME vocabularies, or have defined alternative approaches to exploring a BIBFRAME implementation, such as the NLM work on this topic.³⁷ Although BIBFRAME, Schema.org, BIBFRAME Lite, and other similar standards tend to be at the center of LD discussions for libraries, a number of other standards are emerging that are designed with LD principles in mind. Encoded Archival Description 3 (EAD3), for example, is building in new elements to make better use of Encoded Archival Context—Corporate Bodies, Persons and Families (EAC-CPF) as well as Uniform Resource Identifiers (URIs) from other sources.³⁸ Likewise, a World Wide Web Consortium (W3C) community group has been formed to explore how to extend the Schema.org standard to include better descriptive metadata for digital and physical archives.³⁹

BIBFRAME Lite
<http://bibfra.me>

NLM's efforts to test bibliographic LD schemas as documented in its June 2015 update surfaced test records that followed the BIBFRAME Lite vocabulary where possible, using more granular schemas where necessary.⁴⁰ Per Fallgren's update, the NLM effort largely sought to map BIBFRAME Lite to Resource Description and Access's (RDA) RDF vocabulary, but vocabulary definitions were also drawn from LoC's BIBFRAME vocabulary, MODS RDF, Schema.org, and W3C. One justification offered for this approach is the concern that many efforts are focusing on MARC and BIBFRAME alignment, rather than on designing a vocabulary that is oriented toward a broader range of resources. Alongside these efforts, LoC has continued to advance work on BIBFRAME, launching testing platforms, refining test applications, and contributing to an expansive discussion on BIBFRAME schema issues in the community. The BIBFRAME model has been documented in a series of releases including vocabularies, relationship models, and suggested non-bibliographic applications.⁴¹ Although LoC established a release of BIBFRAME in the summer of 2015, it also continues to refine the standard through a series of proposals.

Outside of the LAM community, LOD has been increasingly adopted to enable better search engine optimization (SEO) and to surface knowledge cards and "rich snippets" in search results and Google's Knowledge Graph.⁴² In 2015, the W3C released a specification for a Linked Data platform that defines a set of systems and system integrations to enable the creation and publication of Linked Data.⁴³ In commercial environments, APIs appear to continue to take precedence over openly published LD. Amazon, for

example, preferences APIs to surface catalog data and enable functional integration. Services such as Alexa (the tool behind Amazon's Echo room system), Marketplace (its tool to publish data on the Amazon catalog), and Mechanical Turk (a system to enable crowd-sourced processing of information) all follow an API over LD model.⁴⁴

Wikipedia: Knowledge Graph

https://en.wikipedia.org/wiki/Knowledge_Graph

Geographic and location-based services including mapping, way finding, and navigation are seeing increasing system integration, but largely through API-based services such as Map APIs, Bluetooth beacon technology, and push-to-mobile interaction techniques. Bluetooth beacons are a good example of the complex relationships that are developing between location-aware services, embedded technology, and the trend toward sensor-based networks, according to Gruman.⁴⁵ These sensors trigger actions in applications based on proximity and can transmit details about the environment, including temperature and time. They can correspondingly log access, provide small bits of information to devices, and help devices triangulate the location of a user in a space by using the proximity information from multiple sensors.

Bluetooth beacons are part of a larger development around the "Internet of Things" (IoT) community in that they can provide description and location information for physical items. Internet-based cameras, Wi-Fi-enabled household products (e.g., televisions, refrigerators, thermostats), and Internet-connected locks and access systems are each contributing to the growing presence of Internet-connected and data-generating devices. As these devices become more common and as their use grows, there is an increasing need to help users bring together these devices and the information they create into a cohesive network that is capable of sharing data as well as inferring new information from shared data among devices.

Ermilov and Auer suggest, for example, that Internet-connected television services could be connected to LD publishers such as DBpedia and IMDb at the client or individual level, enabling a user to actively select content to connect (e.g., a TV guide and IMDb ratings; actor lists and DBpedia entries) from his or her own device rather than working through a centralized service provider that had pre-integrated those services.⁴⁶ At the moment, most IoT technologies work within a specific ecosystem, making it difficult to develop generalized information networks, but some tools, such as Bluetooth beacons, are being designed to work across a range of applications rather than simply within a single application.

LOD Visibility

Within the LAM community, LOD is a commonly discussed topic that tends to have a shared set of values (e.g., make data open, enable reuse, support new uses of data). These values are common in other academic communities, including researchers dedicated to open scholarship and reproducibility as well as creators of data in certain domains. The US government website <http://data.gov>, for example, now provides access to over 150,000 datasets, although in many cases these datasets are serialized in HTML, PDF, and other non-computational document formats. In addition, while the Data.gov site makes items available through a faceted discovery platform, it does not seek to act as an authoritative location for the data and as such does not publish persistent URLs (PURLs). In many cases, however, the data is provided with authorship and license information, two important elements in creating open, if not linked, data.

While LOD is highly visible in the LAM community and is increasingly referenced, by concept if not name, in reproducibility and data publishing communities, it has yet to enjoy widespread understanding or popularization in the press. In fact, searching the web for news stories on Linked Data surfaces more articles from 2000 to 2009, when news companies like the *New York Times* began publishing data as LD, than more recent articles. LD continues to attract funding, however—for example, from the Mellon Foundation, a supporter of the LD4L project; from the Institute for Museum and Library Services (IMLS) in its support for BIBFLOW and the Linked Data for Professional Education programs; and from a range of libraries, archives, and museums that use internal funding to experiment with LD.

New York Times: Linked Open Data (Beta)

<http://data.nytimes.com>

Outside of these funded areas and LAM-focused research threads, whether or not LD and LOD need to enjoy greater visibility in the research community is a topic of debate. Digital humanities programs and communities may be most likely to benefit from LOD experimentation in data publishing as newly published datasets hold the potential to directly drive new threads of research. Likewise, the reproducibility and data science communities could be strong contributors to the evolving practice of LOD in LAM institutions through the development of tools and methods that could be applied to other research domains. The related but as yet unresolved question around visibility is whether or not LD has reached critical mass in the LAM community to ensure further adoption and transformation. The overall lack of visibility of the

role and impact of LD does not help address this issue, although the commitment of large-scale organizations is still heavily influencing how organizations perceive the importance of LD.

Maturity of Vocabularies

The OCLC survey of adoption reviewed earlier in this chapter indicated that LAM institutions are beginning to agree on a series of vocabularies, even if there are areas of ambiguity in how the vocabularies are used or differences of opinion in which vocabularies should be used. One key set of vocabularies that are part of this discussion are BIBFRAME and BIBFRAME Lite and the vocabularies associated with LoC (e.g., Name Authority File, Subject Authority File), as well as the VIAF. The investment in these vocabularies in non-LD formats may ensure that the LD versions enjoy adoption, and in fact they are featured in BIBFRAME and BIBFRAME Lite schemas. How much consensus exists around the higher-level schemas, particularly as framed in the discussion of web visibility, has yet to be seen.

Another important discussion in the LD community centers on the proper fit of vocabularies with different communities of practice. Although BIBFRAME was designed to be a resource-agnostic vocabulary, it has a way to go before it will enjoy broad adoption. As might be expected, the geographic information system (GIS) community has branched out to create its own vocabularies and vocabulary-publishing platform in GeoNames. The discussion around appropriate fit dovetails with related conversations about the perceived value of LD work in general (e.g., how should LAM institutions balance the need for generalized LD models that encourage interoperability with external community members against the need for highly granular internally focused standards)?

Conclusion

Chapter 1 of this issue has served as an overview of the state of LD adoption and sought to catch the reader up from the July 2013 issue of *Library Technology Reports* on Linked Data. This chapter focused in part on the survey completed in 2014 on LD adoption across the LAM community and expanded on identified themes through literature review and exploration of developments in LAM communities.

An original goal of this issue was to gather together the various projects and initiatives underway in the LAM community. As the author engaged in research and studied the results of the 2014 OCLC survey, it became apparent that the LD community has become too large to study comprehensively in a

detailed way. With that in mind, the author is glad to see a revised version of the LD adoption survey being conducted and expects that the results of that survey will be informative for those seeking best practices and guidance on how to launch their own LD projects. Given the fact that the survey results will come shortly after the publication of this issue, it makes sense to focus this work on broad trends and technologies rather than on specific projects and use cases.

In chapters 2 and 3, this issue skims the surface of LD adoption in order to identify representative trends and activities that are currently important in the LD LAM community. Recognizing that these project examples and their importance are situated in the larger context of the web and of the growing use of the Internet of Things and in the broader questions around value and impact, chapter 4 seeks to study the “so what?” questions around LD innovation and adoption.

Notes

1. Karen Smith-Yoshimura, “Linked Data Survey Results 1—Who’s Doing It (Updated),” *hangingtogether.org* (blog), OCLC Research, August 28, 2014, last updated September 4, 2014, <http://hangingtogether.org/?p=4137>; Karen Smith-Yoshimura, “Linked Data Survey Results 2: Examples in Production (Updated),” *hangingtogether.org* (blog), OCLC Research, August 29, 2014, last updated September 4, 2014, <http://hangingtogether.org/?p=4147>; Karen Smith-Yoshimura, “Linked Data Survey Results 3—Why and What Institutions Are Consuming (Updated),” *hangingtogether.org* (blog), OCLC Research, September 1, 2014, last updated September 4, 2014, <http://hangingtogether.org/?p=4155>; Karen Smith-Yoshimura, “Linked Data Survey Results 4—Why and What Institutions Are Publishing (Updated),” *hangingtogether.org* (blog), OCLC Research, September 3, 2014, last updated September 4, 2014, <http://hangingtogether.org/?p=4167>; Karen Smith-Yoshimura, “Linked Data Survey Results 5—Technical Details,” *hangingtogether.org* (blog), OCLC Research, September 5, 2014, <http://hangingtogether.org/?p=4256>; Karen Smith-Yoshimura, “Linked Data Survey Results 6—Advice from the Implementers,” *hangingtogether.org* (blog), OCLC Research, September 8, 2014, <http://hangingtogether.org/?p=4284>.
2. Smith-Yoshimura, “Linked Data Survey Results 1.”
3. Smith-Yoshimura, “Linked Data Survey Results 2”; Karen Smith-Yoshimura, “Results of Linked Data Survey for Implementers,” Excel file, September 5, 2014, OCLC Research, www.oclc.org/content/dam/research/activities/linkedata/oclc-research-linked-data-implementers-survey-2014.xlsx.
4. Smith-Yoshimura, “Linked Data Survey Results 3.”
5. Smith-Yoshimura, “Linked Data Survey Results 4.”
6. Simeon Warner, “LD4L Use Cases,” *Linked Data for Libraries* wiki, last modified by Tom Cramer, May 7, 2015, <http://wiki.duraspace.org/display/ld4l/LD4L+Use+Cases>.

7. Smith-Yoshimura, "Linked Data Survey Results 3."
8. Smith-Yoshimura, "Linked Data Survey Results 4."
9. Smith-Yoshimura, "Linked Data Survey Results 5."
10. Smith-Yoshimura, "Linked Data Survey Results 6."
11. Tim Berners-Lee, "Linked Data," W3C, last updated June 18, 2009, www.w3.org/DesignIssues/LinkedData.html.
12. Nick Crofts, Martin Doerr, and Mika Nyman, "Call for Comments—Linked Open Data Recommendation for Museums," International Council of Museums, accessed July 24, 2015, www.cidoc-crm.org/URIs_and_Linked_Open_Data.html.
13. Smith-Yoshimura, "Linked Data Survey Results 4."
14. Jennifer Marill, "Round Robin Reports—Annual 2015," ALA Connect, June 12, 2015, <http://connect.ala.org/node/240751>.
15. Jim LeBlanc, "Report from Cornell," June 17, 2015, "Round Robin Reports—Annual 2015," ALA Connect, <http://connect.ala.org/node/240751>.
16. Jennifer Marill, "Annual 2015 Report from LC," June 18, 2015, "Round Robin Reports—Annual 2015," ALA Connect, <http://connect.ala.org/node/240751>.
17. Jennifer Marill, "Annual 2015 Report from NLM," June 18, 2015, "Round Robin Reports—Annual 2015," ALA Connect, <http://connect.ala.org/node/240751>.
18. Beth Camden, "Report from Penn," June 19, 2015, "Round Robin Reports—Annual 2015," ALA Connect, <http://connect.ala.org/node/240751>.
19. Bryan Skib, "Report from Michigan," June 23, 2015, "Round Robin Reports—Annual 2015," ALA Connect, <http://connect.ala.org/node/240751>.
20. Ian Hickson, Robin Berjon, Steve Faulkner, Travis Leithead, Erika Doyle, Navara Edward O'Connor, and Silvia Pfeiffer, eds., "HTML5: A Vocabulary and Associated APIs for HTML and XHTML," W3C Recommendation, October 28, 2014, www.w3.org/TR/html5/; Richard Cyganiak, David Wood, and Markus Lanthaler, "RDF 1.1 Concepts and Abstract Syntax," W3C Recommendation, February 25, 2014, www.w3.org/TR/rdf11-concepts.
21. David Wood, ed., "What's New in RDF 1.1," W3C Working Group Note, February 25, 2014, www.w3.org/TR/rdf11-new.
22. Frank Manola, Eric Miller, and Brian McBride, eds., "RDF 1.1 Primer," W3C Working Group Note, February 25, 2014, www.w3.org/TR/2014/NOTE-rdf11-primer-20140225.
23. Erik Wilde, "JSON or RDF? Just Decide," *dretblog*, February 10, 2015, <http://dret.typepad.com/dretblog/2015/02/json-or-rdf-just-decide.html>.
24. Jason A. Clark and Scott W. H. Young, "Building a Better Book in the Browser (Using Semantic Web Technologies and HTML5)," *Code4Lib Journal*, no. 29 (July 15, 2015), <http://journal.code4lib.org/articles/10668>.
25. Suzanna Conrad, "Using Google Tag Manager and Google Analytics to Track DSpace Metadata Fields as Custom Dimensions," *Code4Lib Journal*, no. 27 (January 21, 2015), <http://journal.code4lib.org/articles/10311>.
26. Frank Donnelly, "Processing Government Data: ZIP Codes, Python, and OpenRefine," *Code4Lib Journal*, no. 25 (July 21, 2014), <http://journal.code4lib.org/articles/9652>.
27. Carlo Bianchini and Mirna Willer, "ISBD Resource and Its Description in the Context of the Semantic Web," *Cataloging & Classification Quarterly* 52, no. 8 (2014): 869–87, <http://dx.doi.org/10.1080/01639374.2014.946167>.
28. Gordon Dunsire, "The Role of ISBD in the Linked Data Environment," *Cataloging & Classification Quarterly* 52, no. 8 (2014): 855–68, <http://dx.doi.org/10.1080/01639374.2014.933753>.
29. Jennifer Liss, "DRAFT Checklist for Evaluating Metadata Standards," *Metaware.Buzz* (blog), ALA Metadata Standards Committee, January 20, 2015, <http://metaware.buzz/2015/01/20/draft-checklist-for-evaluating-metadata-standards>.
30. Lucas Mak, Devin Higgins, Aaron Collie, and Shawn Nicholson, "Enabling and Integrating ETD Repositories through Linked Data," *Library Management* 35, no. 4/5 (2014): 284–92, <http://dx.doi.org/10.1108/LM-08-2013-0075>.
31. Péter Király, "Query Translation in Europeana," *Code4Lib Journal*, no. 27 (January 21, 2015), <http://journal.code4lib.org/articles/10285>.
32. Jenny A. Toves and Thomas B. Hickey, "Parsing and Matching Dates in VIAF," *Code4Lib Journal*, no. 26 (October 21, 2014), <http://journal.code4lib.org/articles/9607>.
33. Eric M. Hanson, "A Beginner's Guide to Creating Library Linked Data: Lessons from NCSU's Organization Name Linked Data Project," *Serials Review* 40, no. 4 (2014): 251–28, <http://dx.doi.org/10.1080/00987913.2014.975887>.
34. Terry Reese, "Opening the Door: A First Look at the OCLC WorldCat Metadata API," *Code4Lib Journal*, no. 25 (July 21, 2014), <http://journal.code4lib.org/articles/9863>.
35. Arie Nugraha, "Indexing Bibliographic Database Content Using MariaDB and Sphinx Search Server," *Code4Lib Journal*, no. 25 (July 21, 2014), <http://journal.code4lib.org/articles/9793>.
36. Carol Jean Godby and Ray Denenberg, *Common Ground: Exploring Compatibilities between the Linked Data Models of the Library of Congress and OCLC* (Washington, DC: Library of Congress and Dublin, OH: OCLC Research, January 2015), www.oclc.org/content/dam/research/publications/2015/oclcsearch-loc-linked-data-2015.pdf.
37. Nancy Fallgren, "Experimentation with BIBFRAME at the National Library of Medicine," GitHub, last updated June 24, 2015, <https://github.com/fallgrennj/BIBFRAME-NLM>.
38. "EAD3 Gamma Release," Society of American Archivists, accessed July 24, 2015, www2.archivists.org/groups/technical-subcommittee-on-encoded-archival-description-ead/ead3-gamma-release.
39. Schema Archetypes Community Group, W3C Community and Business Groups, accessed September 12, 2015, <https://www.w3.org/community/archetypes>.
40. Marill, "Annual 2015 Report from NLM"; example records can be found as links on the NLM BIBFRAME

- experimentation report: Fallgren, “Experimentation with BIBFRAME.”
41. Library of Congress, “BIBFRAME Model and Vocabulary,” accessed September 12, 2015, www.loc.gov/bibframe/docs/index.html.
 42. “Rich Snippets (Microdata, Microformats, RDFa, and Data Highlighter),” Google, Webmaster Tools Help, accessed March 11, 2013, <http://support.google.com/webmasters/bin/answer.py?hl=en&answer=99170> (page now discontinued).
 43. Steve Speicher, John Arwe, and Ashok Malhotra, eds., “Linked Data Platform 1.0,” W3C Recommendation, February 26, 2015, www.w3.org/TR/2015/REC-ldp-20150226.
 44. “What’s Available?” Amazon Offerings for Developers, accessed September 12, 2015, <https://developer.amazon.com>.
 45. Galen Gruman, “What You Need to Know about Using Bluetooth Beacons,” *Smart User* (blog), InfoWorld, July 22, 2014, www.infoworld.com/article/2608498/mobile-apps/what-you-need-to-know-about-using-bluetooth-beacons.html.
 46. Timofey Ermilov and Sören Auer, “Enabling Linked Data Access to the Internet of Things,” in *Proceedings: iiWAS2013: 15th International Conference on Information Integration and Web-Based Applications and Services*, ed. Edgar Weippl, Maria Indrawan-Santiago, Matthias Steinbauer, Gabriele Kotsis, and Ismail Khalil, 300–308 (New York: Association for Computing Machinery, 2013), <http://dx.doi.org/10.1145/2539150.2539157>.