IFLA WLIC 2019 Athens

Submitted on: 02.07.2019

From Tension to Support: Leveraging strengths of metadata, context, and prose

John Chapman
Metadata Strategy and Operations, OCLC, Inc., Dublin, Ohio, USA.
E-mail address: chapmanj@oclc.org

Abstract:

In the last year, OCLC has worked with a number of research libraries to determine workflow needs when leveraging Wikibase and its associated services as a descriptive toolset for library collections. Together, we learned important benefits and limitations of this toolset, and insights into the changes it brings to metadata workflows.

Some details of the work of this Linked Data Prototype project have been covered in other presentations, at an operational and/or technical level. In contrast, this paper will explore the lessons learned, within a framework of cataloguing and metadata workflows, and place them in a broader conceptual framework.

The paper will use the custom interface and documentation work that was pursued during the project as evidence for workflow needs not covered by out-of-the-box Wikibase solutions. The paper will illustrate the needs of the metadata professional in using these tools for describing specific types of items in library collections. Throughout, an outline of a co-dependent triad will begin to emerge.

This triad – representing systems that manage content, context, and commentary - will be examined in the light of existing metadata ecosystems. It will also be used to explain the structure and data flows of past and existing metadata efforts. Its limitations, particularly as they vary among material types or other dimensions, will be considered. Finally, it will be used to make some predictions and recommendations for future research and collaboration.

Throughout, the emphasis will be on increasing overall efficiency, reducing rework and latency, and enabling the surfacing of rare and unique materials, while placing them within a broader cultural context.

Keywords: Linked data ; structured data ; data models ; prototypes ; Wikibase

LINKED DATA AND OCLC

OCLC has been creating and publishing linked data for many years. It is the technical infrastructure provider for the Virtual Internet Authority File (VIAF),[1] one of the most heavily-used library linked data resources. In addition, OCLC has created and maintains Faceted Access to Subject Terminology (FAST),[2] which provides a faceted vocabulary of terms that is based on the Library of Congress Subject Headings (LCSH).[3] More recently, OCLC has created Works and Person entities through clustering diverse sources of data including bibliographic and authority data.[4] Building from this foundation were multiple projects, pilots, and prototypes; among them, the Person Entity Lookup Service,[5] EntityJS and other discovery tools,[6] and the Metadata Refinery.[7] In 2017, OCLC management created an internal Linked Data Task Force, charging it with pulling together lessons learned from past projects - and exploring future opportunities.

THE LINKED DATA PROTOTYPE

As the Linked Data Task Force concluded its research, OCLC management asked them to deliver a project, within approximately nine months, that would deliver something that would advance the discussion of linked data in the marketplace and could be turned into a basis for future products and services. The resulting project, the OCLC Linked Data Prototype, known colloquially as "Project Passage," ran from January to September 2018; a lengthy analysis is found in an OCLC research report.[8]

The tight schedule for the effort meant that we needed to focus on constrained deliverables:

> 1: **Develop an entity ecosystem** that facilitated: a) Creation and editing of new entity descriptions; b) Connecting entities to the Web.
>
> 2: **Build a community of users** who could: a) Create and curate data in the ecosystem, and b) Imagine and propose workflow uses.
>
> 3: **Provide services** to a) Reconcile data, b) Explore the data.[9]

---

[1] VIAF: Virtual Internet Authority File. https://viaf.org
[2] SearchFast. https://fast.oclc.org/searchfast/
[3] Library of Congress Subject Headings - LC Linked Data Service: Authorities and Identities. http://id.loc.gov/authorities/subjects.html
[4] OCLC, Inc. "Data strategy [WorldCat]." OCLC, Inc. (Dublin, Ohio, USA): 2019. https://www.oclc.org/en/worldcat/data-strategy.html
[5] "OCLC to launch linked data pilot with seven leading libraries." OCLC, Inc. (Dublin, Ohio, USA): 11 September 2015. https://cdm15003.contentdm.oclc.org/digital/collection/p15003coll6/id/382/rec/169
[6] Washburn, Bruce, and Jeff Mixter. "Looking inside the library knowledge vault." OCLC Research Library Partners, Works in Progress Series: 12 August 2015. https://www.oclc.org/content/dam/research/events/2015/library-knowledge-vault-webinar-2015.pptx
[7] Mixter, Jeff. "Linking your Data." Canadian Linked Data Summit (Toronto, Canada), 24 October 2015. https://www.mcgill.ca/clds/files/clds/mixter-en.pptx
[8] Godby, Jean; et al. "Creating Library Linked Data with Wikibase: Lessons Learned from Project Passage." OCLC Research, forthcoming. [2019]

[9] Pace, Andrew K., and Jason Kovari. "Prototyping a Linked Data Platform for Production Cataloging Workflows." CNI Spring 2018 Membership Meeting (San Diego, California): 13 April 2018. https://www.oclc.org/content/dam/research/presentations/pace/2018-CNI-Kovari-Pace.pptx

From the start, we sought to engage the user community in creating and editing the data. The deliverables also included the intentional cultivation and maintenance of a community of users, to help evolve the ecosystem itself, and help define its future potential. We would also provide services for the community to connect local data to the ecosystem we hosted, and we would provide interfaces that would help show its potential for discovery.

The technical team settled on the approach of using Wikibase. They have presented in multiple venues on the technology and the reasons for its selection.[10] Due to the similarity with Wikidata's user interface, the lessons learned about end-user behavior with the Wikibase software are broadly applicable to the public Wikidata interface.

The entity ecosystem we built, which comprised Wikibase and locally-developed data processing tools, started with approximately 1 million entities. High-quality and heavily cross-referenced, they were based on the overlap of Wikidata, VIAF resources, and the Library of Congress Name Authority File (LCNAF). Later on, approximately 40,000 more Work entities were added.

Workflow questions

Many library linked data projects start with mass conversion of legacy data. OCLC is a major data aggregator with significant experience creating, transforming, and distributing linked data. However, underlying the purpose and utility of WorldCat is the shared creation and maintenance of its data. Conversion and clustering, as we saw with the Works and Persons data we generated, didn't deliver the quick feedback libraries were used to seeing. When cataloguers made changes to bibliographic records in WorldCat, or in their local catalogue, they could switch to their discovery interface and see those reflected immediately.

Additionally, in our projects that included archivists and special collections librarians, we learned that, regularly, the people or places their materials described weren't represented in library catalogues or authority files. Naturally, they also were absent from entity aggregations we had built from those databases. They asked us to create tools to directly add new entities.

These user needs made it clear that this wasn't strictly a metadata creation and publishing experiment, but instead an exploration of ways to improve the fit of linked data to discovery functionality, and to improve the efficiency and quality of metadata work.

Some of the issues around discovery are plain. Quality of linked data, as it is represented by well-modeled, predictable, and consistent linking to well-maintained outside resources, was of paramount importance. But outside of the "usual suspects" – the name and subject authority files and bibliographic records that commonly link to them – we knew there was significant interest, but not enough experience, in leveraging these other linked data to provide context. How should this context be reflected to the user?

Other workflows were centred more on the traditional "cataloguer" use cases: mechanics of creating or finding accurate information and adding it to the database. How could this be made more efficient? What if new vocabularies are desired? What does "copy cataloguing"

[10] Washburn, Bruce, and Carl Stahmer. Prototyping a Linked Data Platform for Production Cataloging Workflows." LD4 Workshop 2018 (Palo Alto, California): 2 May 2018. Available at: https://www.oclc.org/content/dam/research/presentations/washburn/LD4-Washburn-Stahmer.pptx

look like? Is it even needed? How much needs to be added to the database, and how much can be simply linked to, in order to leverage or query at the time of need?

Gaps - and new functionality to fill them

The first activity for project participants that involved editing was to create interrelated entities based on translations of works. This was meant to illuminate translation-modeling work previously done by OCLC Research[11], but also to test what we predicted were powerful but unconstrained ways to represent relationships between resources, names, and labels crossing multiple languages and alphabets. It allowed users to use the so-called "Fingerprint data." The fingerprint data is generally used to provide brief informational labels for the purposes of disambiguation, and to provide this information in a number of relevant languages and/or alphabets so that it can be found through string-based searching.

As the users moved to edit existing entities and create new ones, we noticed a few important gaps. The first was an end-user discovery application. We needed a tool that displayed the structured data in a friendlier mode. OCLC staff built from their experience with entity display - especially the EntityJS work – and built a tool called the "Explorer."[12]

Crucially, this tool did not simply display the structured data and labels from the fingerprint data. Instead, it adapted specific URIs from the structured data, representing identifiers for external databases, to pull images and text descriptions. These images and associated descriptive text were pulled in and incorporated in the web page only when the web page was built.

The Explorer also used basic queries against the entity ecosystem to determine which entities linked to the one currently being viewed. These assertions were then collected by type and displayed alongside the entity data. As an example, the Explorer view for Barack Obama would list the books for which he is the author, as well as works that have him as a topic, though these statements are stored in entities other than the one being viewed.

Key discoveries

When the participants began entering information in the Fingerprint data, we got our first hints of a core issue in the system, and the first hints at the theme of this paper. Due to the importance of this data for searching, and to the lack of free text fields, this data was sometimes filled with long descriptions of the material. While well-known person entities might have the informative Wikipedia context we could display in Explorer, many creative works did not.

It was readily apparent this was not ideal, but the default Wikibase interface does not provide a place for non-structured descriptive text for the resource. Despite the power of structured data, "interpretive context is still needed for most types of resources, especially for digital images and archival collections."[13] But adding it to Fingerprint data overloaded a field that was meant for short names, titles, and other similar labels. The interface and tools made it

[11] OCLC Research. "Multilingual Bibliographic Structure." OCLC, Inc. (Dublin, Ohio, USA): 2019. https://www.oclc.org/research/themes/data-science/multilingual-bib-structure.html
[12] Godby, et al,; see links to Explorer views on p 39-40.
[13] Godby, et al.; p 61.

clear this was not the intended purpose, with the difference in functionality and visual appearance of this textual display.

But in many cases, the alternative was daunting: connecting a complex network of relationships to a large set of entities, in some cases creating entirely new entities in order to support these assertions. With each new entity, the user saw the possibilities multiplying: to put a photograph in context, a photographer, a location, a subject, a collection might all need to be added to the ecosystem; for each of those entities, several more entities might follow.

OCLC engineers built a tool called the "Retriever" to streamline the work.[14] It was an interface set up to query specific external data sources, for example Wikidata, and then construct new entities in the prototype system based on – and already linked to – the source entity. The Retriever allowed the user to preview the labels and assertions that were automatically detected, and also confirm their accuracy; they could be deleted or modified before the new entity was added to the ecosystem.

How far to go?

With the new tool in place, a question that had been asked early on in the prototype gained new importance: How far should users go in adding entities to this entity ecosystem? The Wikibase platform, with the addition of tools like the Retriever, is powerful enough that we may be tempted to overuse it:

> Expert communities will need to determine where to redraw the line between structured data and text, now that a software solution such as the Wikibase platform makes it easier to create structured data than is possible with current standards and tools.[15]

After the structured data describing these new entities are added, even with new properties extending the ontology, and using the powerful Explorer application to view the results, the aggregate results were still missing something.

> …it was not obvious what should be expressed as structured data and how this information could be used in a discovery interface to build an interpretation that is superior to the products of existing practice…. human-readable narrative or expository text will continue to serve a useful function…. [16]

The Passage report cites the experience of Temple University, who saw the complexity of the structured data as a potential trap:

> …important structured statements in a knowledge graph may be many steps away from the entity for the primary item of interest. A mediated view or query must be designed to reveal this deeply buried relationship,

[14] Washburn, Bruce. "OCLC Project Passage User Interface: Assisting the editing workflow." 2019 LD4 Conference (Boston, Massachusetts, USA): 11 May 2019. https://drive.google.com/open?id=1XAdHZZJsyWBPaDJNrQIDpwNGSiARrOtq

[15] Godby, et al.; p 55.
[16] Godby et al.; p 54-55.

> which seems more complex and less trustworthy than a human-readable page….[17]

As we saw earlier, the default Wikibase paradigm does not support interpretive context; the fingerprint data is a tempting but dysfunctional way to convey such narrative commentary. But it is clearly needed.

> Both structured and narrative data are still necessary…. Perhaps this realization is also reflected in the configuration of the Wikimedia ecosystem. There, Wikidata is…a graph of structured data featuring factual, verified statements about items and properties. Wikipedia is a repository for human-readable exposition…. And Wikimedia Commons contains structured data about digital media, including rights statements and technical metadata.[18]

This triad – Wikidata, Wikipedia, and Wikimedia Commons – suggests a 'separation of concerns' that sheds important light on the possibilities and limitations of traditional library systems – and the library metadata ecosystem itself – as they continue to evolve to meet the demands of web-based discovery.

A THREE-PART DIVISION

The questions that kept arising around "how far to go," and how Wikibase might connect to other pools of library data, including MARC catalogues, led us to some insights. Prominent among these was a functional model of web-based library data tools.

At its highest level, the model talks about Context, Commentary, and Content. Broadly speaking, **Context** consists of formal statements; **Commentary** is prose or annotation; and **Content** comprises information objects and technical metadata about their identification, maintenance and provision.

I worked up a quick diagram, with three functional areas in three circles. This sketch quickly became known as "Passage Bubbles." I quickly extended it to cover a few more ideas.

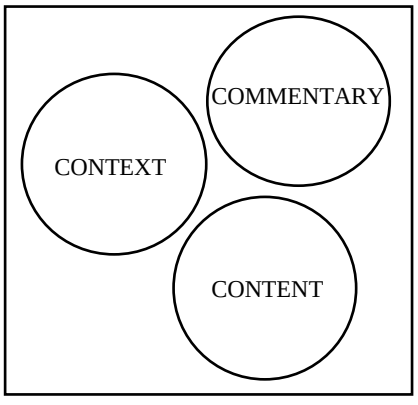

Figure 1. Originally labelled "what type of information does each part focus on?"

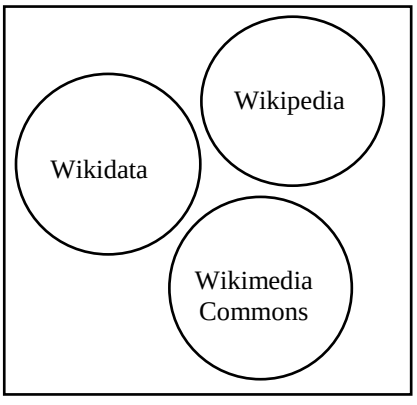

Figure 2. Originally labelled "what are the corresponding systems in wiki-ville?"

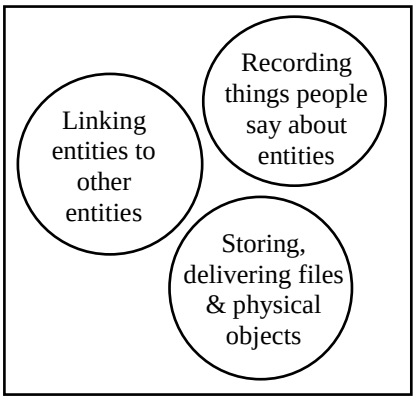

Figure 3. Originally labelled "what roles do the different systems play?"

[17] Godby et al, p 61.
[18] Godby et al, p 45.

These diagrams do not show the interaction between the spheres, or "Bubbles." In our discussions, we assumed these would be assertions or statements, such as semantic triples. We also prepared some additional Bubble diagrams. One maps different library systems:

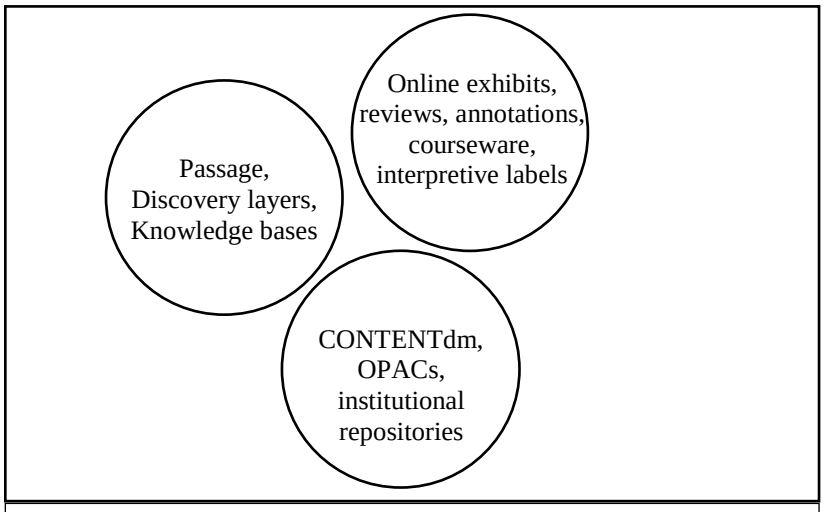

Figure 4. Originally labelled: "what are the systems in library-land?"

Notable here is the separation of "OPACs" and Passage. This summarizes our assumption that the purpose of tools like Wikidata is to provide a contextual "layer." This is something to which MARC-based systems, or other record-based systems including repositories or digital asset management systems, can refer.

I also prepared additional Bubbles diagrams for user and agent types, and mapping to the "User Tasks" of the Functional Requirements for Bibliographic Records, or FRBR.[19] These are more complex, as they show multiple overlaps. For clarity, the entire set of diagrams is available online.[20]

How well does a MARC-based system fit?

Starting from its history as a network-capable improvement on physical card-based tools for inventory control and local discovery, the MARC bibliographic record is relied upon for a broad range of purposes. When this range of purposes is embodied in a single document, strains begin to appear.

The MARC bibliographic record's functions are stretched across all three Bubbles. Context can be seen in its references to controlled vocabularies, authority records, and formal links to other resources in the catalogue. Content is seen in its relationship to the Item or holdings record, the physical description fields, and notes that describe structure. Commentary is the least obvious, but this in itself points to difficulties.

[19] Tillett, Barbara. "What is FRBR? A conceptual model for the bibliographic universe." Library of Congress Cataloging Distribution Service (Washington D.C.): 2004. https://www.loc.gov/cds/downloads/FRBR.PDF

[20] Chapman, John. "Passage Bubbles." OCLC, Inc. (Dublin, Ohio, USA): June 2018. http://bit.ly/PassageBubbles

Commentary in MARC records takes several forms, often as notes. They use unstructured prose to convey a cataloguer's professional judgment on the material being described, how it was obtained, or how to make use of it, and by whom this is ideally done.

Not every free-text field is commentary; for example, title subfields (when they are not defined by title authorities) are analogous to the Fingerprint data of the Wikibase record: labels for the purposes of creating title search access, or for placing things in a sequence or structure.

Any field in a MARC bibliographic record must be considered as a part of a whole. The accuracy of the record itself is dependent on validation routines, structural and syntactic, that take into account the entire record. Even those fields that are solitary, not defined by or defining other fields, lose their meaning separated from their container record and its identifier. Much of the semantics of the record, for example material type, is dependent on the complex interplay of a number of different values found in separate parts of the record.

MARC is likely to have a future as its original design intended – as an exchange format for Content-related metadata, and will likely continue to play a strong role for many libraries in acquisitions and circulation tasks, even as they adopt linked data for discovery. This would support the idea that its adaptation for contextual discovery, and its relatively closed nature, was not well adapted for web-based discovery.

The Bubbles also provide an insight into the provision of reviews, commentary, and annotation in MARC-based systems. There is a glaring mismatch between the locked-down nature of the MARC record - on which so much depends – and the expectations of the public for agile correction, addition, and reviews. Handling both the "privileged commentary" of a MARC record's notes, and information supplied by users, displays the incomplete modelling of commentary within the MARC record. It also suggests the potential for pulling these references, no matter who creates them, to other commentary systems that can manage identities and rights more robustly.

Wikimedia's fit

We discovered later that Wikimedia volunteers had arrived at some of the same conclusions as we had about the roles of the Wikimedia tools. Andrew Lih of Wikimedia DC characterized[21] the different roles of the Wikimedia space:

Wikimedia Commons	*Upload, editing, annotation, categorization*
Wikipedia	*Summarizing, referencing, neutral point of view*
Wikidata	*Querying, multimedia, visualization, tool-building*

Lih's list shows some real-life differences from our theoretical Bubbles model: Wikimedia Commons incorporates some categorization, which blurs the border with categories and prose description in Wikipedia, and with links to formal vocabularies in Wikidata, for example genre or material type.

[21] Lih, Andrew. "Beyond Wikipedia: interconnecting human knowledge through Wikidata." ALA 2019 Annual Conference (Washington, DC, USA): 22 June 2019. http://bit.ly/ala19wikidata

A useful discussion of the role of categories in Wikimedia Commons can be found in an interview with George Oates.[22] Oates discusses the shortcomings of the current state of categories in Wikimedia Commons. Oates discusses the extremely unorganized set of categories and the fact that they include both technical categories, such as image resolution, as well as categories depending on the semantics or content of the things depicted or described. While category-based discovery could certainly live in systems devoted to Commentary or Context, having at least some Categories in the Content zone is clearly necessary and useful. It is also clear that currently, the Wikimedia Commons is not always consistent in distinguishing metadata about a real-life object and metadata about a digital file describing it. However, there is great potential in the Commons – and its large user base – continuing to explore its impact on users.

In addition, the large Wikicite[23] project is suggestive of models of Commentary that would include scholarly communication and citation.

QUESTIONS FOR THE FUTURE

What's missing from the Bubbles model?

One of the most significant issues regarding distributed metadata is the notion of discovery. Specifically, when resources are contextualized using identifiers, labels must be provided for the purposes of text-based searching and human-readable display. While much of the power and promise of distributed, linked open data is in the ability to customize and localize these labels, there is much yet to be determined in how search interfaces might be able to leverage remote but linked labels and context.

One clear engineering solution to this issue of uneven and unpredictable on-demand searching is caching. Normally, systems like to periodically and systematically retrieve and cache information that is unlikely to change often. For the more dynamic sets, system designers will need to consider more complex models of notification, subscription, mirroring, or distribution. But the library community will also need to come to political and operational agreements on centralization and decentralization, which will further impact redundancy, preservation, and other concerns. The result may be a blurring or overlap of the Bubbles.

What about representing local collections in heavily interconnected and interdependent systems? There are many questions, political and operational. How will metadata important or meaningful only to local audiences be stored? How can it be made more widely understood? What about situations in which the metadata is private, confidential, or puts people at risk? A distributed model that leverages local metadata that uses globally (or regionally) defined context must allow for one-way references from the local to the global, while responding to requests for global context in a manner that can preserve local anonymity.

[22] Vershbow, Ben, and Sandra Fauconnier. "How could Wikimedia Commons be improved? A conversation with designer George Oates." Wikimedia Foundation: 29 October 2018. https://wikimediafoundation.org/2018/10/29/george-oates-conversation/

[23] Wikicite. [Viewed 10 June 2019] https://meta.wikimedia.org/wiki/WikiCite

What's next?

Currently, multiple library-related projects are using Wikidata and other related technologies. At the national library level, BnF and ABES recently issued a request for help setting up a Wikibase instance for authority file creation.[24] Other national libraries are currently experimenting with publishing data using Wikidata[25], and their discoveries will have large impacts on the future of contextual discovery and metadata creation. Late in 2018, Stanford University hired a "Wikimedian-in-Residence" to work out the implications of Wikidata for the Linked Data For Production (LD4P) grant funded by the Mellon Foundation.[26] The momentum in Wikidata experimentation was recently increased by the Association of Research Libraries' white paper on Wikidata[27], which provides crucial information for libraries seeking to experiment further.

OCLC is further exploring the concepts around the Bubbles, by further exploring Wikibase's ability to place special collections and archives in Context; tying together the extensive Commentary created by users of the CONTENTdm platform; and that software's successful implementation of the IIIF standard for providing Content.

Acknowledgments

The author thanks the library participants in the Linked Data Prototype; and his colleagues at OCLC, namely Sara DeSmidt, Jean Godby, Jeff Mixter, Stephan Schindehette, Taylor Surface, Bruce Washburn, and Jeff Young. Their dedication and hard work were fundamental to the success of the Linked Data Prototype. Any errors or mischaracterizations are the sole responsibility of its author.

[24]Bibliothèque nationale de France [BnF] and Agence bibliographique de l'enseignement supérieur [ABES]. "Appel d'offre pour une preuve de concept : WikiBase comme infrastructure pour le Fichier national d'entités." BnF (Paris, France): 3 April 2019. https://www.transition-bibliographique.fr/2019-04-02-appel-offre-preuve-concept-wikibase-pour-fne/

[25] Miller, Matt. "Wikidata at Library of Congress." ALA Annual Conference 2019 (Washington, DC), 22 June 2019. http://bit.ly/ala19wikiloc

[26] Futornick, Michelle. "Wikimedian-in-Residence position at Stanford University." LD4P2 Linked Data for Production: Pathway to Implementation: 27 Aug 2018. https://wiki.duraspace.org/display/LD4P2/Wikimedian-in-Residence+position+at+Stanford+University

[27] Association of Research Libraries. "ARL White Paper on Wikidata" (April 2019). https://www.arl.org/resources/arl-whitepaper-on-wikidata/

References

Association of Research Libraries. "ARL White Paper on Wikidata" (April 2019). https://www.arl.org/resources/arl-whitepaper-on-wikidata/

Bibliothèque nationale de France [BnF], and Agence bibliographique de l'enseignement supérieur [ABES]. "Appel d'offre pour une preuve de concept : WikiBase comme infrastructure pour le Fichier national d'entités." BnF/ABES (Paris, France): 3 April 2019. https://www.transition-bibliographique.fr/2019-04-02-appel-offre-preuve-concept-wikibase-pour-fne/

Chapman, John. "Passage Bubbles." OCLC, Inc. (Dublin, Ohio, USA): June 2018. http://bit.ly/PassageBubbles

Futornick, Michelle. "Wikimedian-in-Residence position at Stanford University." LD4P2 Linked Data for Production: Pathway to Implementation: 27 Aug 2018. https://wiki.duraspace.org/display/LD4P2/Wikimedian-in-Residence+position+at+Stanford+University

Godby, Jean; et al. "Creating Library Linked Data with Wikibase: Lessons Learned from Project Passage." OCLC Inc. (Dublin, Ohio USA): forthcoming.

Library of Congress Subject Headings - LC Linked Data Service: Authorities and Identities. [Viewed 10 June 2019] http://id.loc.gov/authorities/subjects.html

Lih, Andrew. "Beyond Wikipedia: interconnecting human knowledge through Wikidata." ALA 2019 Annual Conference (Washington, DC, USA): 22 June 2019. http://bit.ly/ala19wikidata

Miller, Matt. "Wikidata at Library of Congress." ALA Annual Conference 2019 (Washington, DC, USA): 22 June 2019. http://bit.ly/ala19wikiloc

Mixter, Jeff. "Linking your Data." Canadian Linked Data Summit (Toronto, Canada): 24 October 2015. https://www.mcgill.ca/clds/files/clds/mixter-en.pptx

OCLC, Inc. [Viewed 10 June 2019] https://www.oclc.org/

OCLC, Inc. "Data strategy [WorldCat]." OCLC, Inc. (Dublin, Ohio, USA): 2019. [Viewed 10 June 2019] https://www.oclc.org/en/worldcat/data-strategy.html

OCLC Research. "Multilingual Bibliographic Structure." OCLC, Inc. (Dublin, Ohio, USA): 2019. [Viewed 10 June 2019] https://www.oclc.org/research/themes/data-science/multilingual-bib-structure.html

OCLC, Inc. "OCLC to launch linked data pilot with seven leading libraries." OCLC, Inc. (Dublin, Ohio, USA): 11 September 2015. https://cdm15003.contentdm.oclc.org/digital/collection/p15003coll6/id/382/rec/169

Pace, Andrew K., and Jason Kovari."Prototyping a Linked Data Platform for Production Cataloging Workflows." CNI Spring 2018 Membership Meeting (San Diego, California,

USA): 13 April 2018. https://www.oclc.org/content/dam/research/presentations/pace/2018-CNI-Kovari-Pace.pptx

SearchFast. [Viewed 10 June 2019] https://fast.oclc.org/searchfast/

Tillett, Barbara. "What is FRBR? A conceptual model for the bibliographic universe." Library of Congress Cataloging Distribution Service (Washington D.C.): 2004. https://www.loc.gov/cds/downloads/FRBR.PDF

Vershbow, Ben, and Sandra Fauconnier. "How could Wikimedia Commons be improved? A conversation with designer George Oates." Wikimedia Foundation: 29 October 2018. https://wikimediafoundation.org/2018/10/29/george-oates-conversation/

VIAF: Virtual Internet Authority File. [Viewed 10 June 2019] https://viaf.org

Washburn, Bruce, and Jeff Mixter. “Looking inside the library knowledge vault.” OCLC Research Library Partners, Works in Progress Series (12 August 2015). https://www.oclc.org/content/dam/research/events/2015/library-knowledge-vault-webinar-2015.pptx

Washburn, Bruce, and Carl Stahmer. "Prototyping a Linked Data Platform for Production Cataloging Workflows." LD4 Workshop 2018 (Palo Alto, California, USA): 2 May 2018. Available at: https://www.oclc.org/content/dam/research/presentations/washburn/LD4-Washburn-Stahmer.pptx

Wikicite. [Viewed 10 June 2019] https://meta.wikimedia.org/wiki/WikiCite