

ResourceSync: Leveraging Sitemaps for Resource Synchronization

Bernhard Haslhofer
Cornell University, Information
Science
bh392@cornell.edu

Martin Klein
Los Alamos National
Laboratory
mklein@lanl.gov

Simeon Warner
Cornell University Library
simeon.warner@cornell.edu

Robert Sanderson
Los Alamos National
Laboratory
rsanderson@lanl.gov

Herbert Van de Sompel
Los Alamos National
Laboratory
herbertv@lanl.gov

Carl Lagoze
University of Michigan, School
of Information
clagoze@umich.edu

Michael L. Nelson
Old Dominion University
mln@cs.odu.edu

ABSTRACT

Many applications need up-to-date copies of collections of changing Web resources. Such synchronization is currently achieved using ad-hoc or proprietary solutions. We propose ResourceSync, a general Web resource synchronization protocol that leverages XML Sitemaps. It provides a set of capabilities that can be combined in a modular manner to meet local or community requirements. We report on work to implement this protocol for arXiv.org and also provide an experimental prototype for the English Wikipedia as well as a client API.

Categories and Subject Descriptors

H.4 [Information Systems Applications]: Miscellaneous

Keywords

Web, Resource Synchronization, ResourceSync

1. INTRODUCTION

Synchronization of resources from one Web-based system, a *source*, to another, a *destination*, is frequently important. It may be necessary to ensure reliable access to a set of resources, to provide backup copies for preservation purposes, or to leverage computational resources or tools available at one server but not another. We consider three examples of synchronization used by popular services:

- The arXiv.org collection of physics, mathematics, and computer science articles that exist on a primary server but are mirrored at other servers worldwide.
- Structured Web data sources such as DBPedia that are synchronized with changes in their unstructured counterparts (Wikipedia).

- The data.europeana.eu service, which aggregates meta-data from many remote sources across Europe.

The synchronization problem is well-known and various solutions are available. One can use *rsync* [5] to synchronize local and remote file systems, use OAI-PMH¹ to synchronize metadata between repositories, synchronize files between two machines via WebDAV², or install Dropbox or Google drive to synchronize local files or documents with cloud services. However, all of these mechanisms are proprietary or ad-hoc solutions, which position the important function of synchronization as an outlier among other functionality that has been comfortably incorporated with the Web Architecture (e.g., REST and Linked Data).

We propose ResourceSync [6, 3] as a framework for a Web-based synchronization mechanism, which leverages the widespread adoption of XML Sitemaps³. It provides a modular set of synchronization components for baseline synchronization, incremental synchronization, and pull- or push-based change awareness that can flexibly be combined to meet a variety of synchronization requirements. Our presentation will focus on the following issues:

1. We introduce the ResourceSync framework, which is now available in a first beta draft (<http://www.openarchives.org/rs/0.5/>). In the simplest case, it requires only that a source exposes an XML Sitemap listing resources with last modification information.
2. We report on the experiences from implementing ResourceSync for arXiv.org and the English Wikipedia and discuss how it might support information aggregation from many diverse sources in Europeana.

All ResourceSync implementations, including a Python client library and a simulator for testing purposes, are available in a Git repository: <https://github.com/resync>.

¹<http://www.openarchives.org/pmh/>

²<http://tools.ietf.org/html/rfc4918>

³<http://www.sitemaps.org/>

2. SYNCHRONIZATION SCENARIOS

2.1 arXiv.org

arXiv (<http://arXiv.org>)⁴ is a well-known and heavily used repository of scholarly articles in physics, mathematics, computer science and related disciplines. It has over 800,000 articles with an average of about 1.5 revisions per article. For each version there is a separate metadata record and the full-text package (PDF, tar.gz, etc.) giving a total of approximately 2.4 million resources. New articles and revisions are added at the rate of about 75,000 per year, and there are also occasional metadata changes such as the addition of bibliographic information for journal versions of articles. Updates are made public at 8pm eastern US time each day Sunday through Thursday, with an average of 1,800 resources updated at that moment at that time.

We consider two synchronization use cases: the first is the synchronization of content to mirror sites⁵ which are under direct arXiv control, and the second is synchronization of content to other independent services⁶ or to researchers for bibliometric and scientometric analysis. In the first case, the goals are high consistency, moderate latency, and robustness to global network outages. There is also the need for the system to automatically recover from outages without human intervention. In the second case, the goal is to make resource and update information publicly available so that any other service may synchronize at the frequency it needs without the need for any out-of-band communication.

The current mirroring system uses a process of an HTTP trigger from the main site, an HTTP pull of a list of changed objects specific to the particular mirror site, HTTP download of the resources listed, an HTTP request to verify when the mirror has completed downloading, and then verification (via HTTP HEAD) by the main site which updates the list of out-of-sync items for the particular mirror. The process is periodically repeated as long as there are updates for that mirror. This system is limited to a trusted set of servers operating with the same internal organization and is not available publicly. It also does not support an audit process to check synchronization so rsync is used periodically. Switching to a standardized resource-centric framework could combine efficient updates, the ability to do periodic audits, public synchronization capability, and reduce the burden of maintaining an ad-hoc system.

2.2 wikipedia.org

The multilingual online encyclopedia Wikipedia currently contains over 23 million articles in various languages and is maintained by about 100,000 active contributors. There are 285 language editions: the English Wikipedia with 4 million articles is the largest one, followed by the German, French, and Dutch editions, which all have between 1 and 4 million articles⁷. According to measurements reported in [2], about 1.4 article pages are updated each second on Wikipedia which amounts to 120,000 page updates per day.

⁴The arXiv statistics come from Simeon Warner, one of the authors of this paper.

⁵<http://arXiv.org/help/mirrors>

⁶Example services include the UK Institute of Physics <http://eprintWeb.org/> site, or the Math Front at UC Davis <http://front.math.ucdavis.edu/>

⁷<http://en.wikipedia.org/wiki/Wikipedia>

We envision two main synchronization use cases for Wikipedia: first there are large structured knowledge bases such as Freebase and DBpedia which are increasingly used in combination with information retrieval techniques (e.g., DBpedia Spotlight [4], Google Knowledge Graph⁸). They heavily reuse or even mirror information from Wikipedia and therefore need to refresh this information when remote resources change. Second, there are publishers and media agencies like the New York Times or the BBC, which link entities in their information space with entities in Wikipedia and reuse information (e.g., article abstracts) in their own information space. Keeping these information sources in sync is an important component of sustaining the timeliness of the news source.

Currently Wikipedia exposes article metadata and change information via a non-public OAI-PMH endpoint, and some editions also push article change information via a dedicated IRC channel. This means that there is no uniform, Web-based solution that allows clients to replicate and periodically synchronize information from Wikipedia.

2.3 data.europeana.eu

Europeana provides access to more than 20 million books, paintings, films, museum objects and archival records that have been digitized throughout Europe, gathered from hundreds of individual institutions, with the help of dozens of data aggregators and providers. The initial Linked Open Data release contains metadata on approximately 2.4 million texts, images, videos and sounds. These collections encompass more than 200 cultural institutions from 15 countries. While the 10 largest data providers contribute 80% of all data, the remaining 20% are contributed by smaller institutions. Two data providers even contribute only one single object to the current dataset [1].

In Europeana, Web-based resource synchronization can serve two purposes: first, Europeana internally needs to periodically aggregate metadata from its data providers. Since Europeana provides an entry point for search and retrieval over aggregated objects, it is in the data provider's interest that Europeana has a consistent view over these objects. Second, Europeana could also provide a synchronization endpoint for external services that consume and make use of data provided by Europeana.

At the moment, the Europeana-internal metadata aggregation mechanism is based on OAI-PMH and manual data transfer from the data providers to Europeana, possibly via intermediate aggregators. External data consumers can download data or use available Europeana data dumps but have no means to synchronize resources.

3. RESOURCESYNC

In order to allow a destination to initially synchronize with a source, the synchronization process at the destination must be able to retrieve a list of source resources for which synchronization is intended — we denote this process as *baseline synchronization*. Subsequently it can perform *incremental synchronization* to keep its copies in-sync with the corresponding resources at the source. Finally, the destination can perform an *audit* to check that its copies match the corresponding resources at the source. The central frame-

⁸<http://www.google.com/insidesearch/features/search/knowledge.html>

work components enabling these processes are *resource lists* and *change lists*. ResourceSync also supports the notion of dumps for packaging resources, cross-linkage of related resources, patching resource representations, etc. Further information on these and other capabilities is given in the ResourceSync beta draft.

3.1 Resource List

Baseline synchronization requires that a source exposes the list of resource URIs it conveys for synchronization. In its most basic form (see Listing 1), a resource list is an XML Sitemap with an additional element expressing that the Sitemap implements resource list capability. Since a resource list presents a snapshot of a Source’s resources at a particular point in time, this element also carries the date-time of the resource list’s most recent update.

```
<?xml version="1.0" encoding="UTF-8"?>
<urlset xmlns="..."
  xmlns:rs="...">
  <rs:md capability="resourcelist"
    modified="2013-01-03T09:00:00Z"/>
  <url>
    <loc>http://example.com/res1</loc>
  </url>
  <url>
    <loc>http://example.com/res2</loc>
  </url>
</urlset>
```

Listing 1: Resource List

3.2 Change List

Incremental synchronization is an optimization over baseline synchronization. If supported by both source and destination, it can reduce latency caused by the transfer of possibly large resource lists. Instead of retrieving the list of available resources, destinations can retrieve atomic resource state change information bundled in change lists (see Listing 2). Change lists are also expressed as Sitemaps.

```
<?xml version="1.0" encoding="UTF-8"?>
<urlset xmlns="..."
  xmlns:rs="...">
  <rs:md capability="changelist"
    modified="2013-01-03T11:00:00Z"/>
  <url>
    <loc>http://example.com/res2.pdf</loc>
    <lastmod>2013-01-02T13:00:00Z</lastmod>
    <rs:md change="updated"/>
  </url>
  <url>
    <loc>http://example.com/res3.tiff</loc>
    <lastmod>2013-01-02T18:00:00Z</lastmod>
    <rs:md change="deleted"/>
  </url>
</urlset>
```

Listing 2: Change List

Both resource lists and change lists can carry additional meta-information (e.g., hash, content length, mime-type) to facilitate resource synchronization and audit.

4. DEMOS

We have implemented ResourceSync prototypes for two use cases: arXiv.org⁹ and the English Wikipedia¹⁰. Each prototype exposes a resource list and a change list containing information about resource changes. Resource lists are written by a periodic process, which is triggered daily in the case of arXiv and by the availability of a new dump in the case of Wikipedia. The arXiv prototype writes one change list file per day, whereas the Wikipedia prototype keeps the most recent changes in memory and periodically writes them to a file. In the case of Wikipedia changes are recorded from a dedicated IRC channel (#en.wikipedia), in the case of arXiv the changes are visible in the arXiv database and periodically written out by a batch process. We are also investigating the possibility of implementing ResourceSync for Europeana.

5. SUMMARY

We propose ResourceSync, as a general Web resource synchronization protocol that leverages XML Sitemaps. It provides a set of capabilities that can be combined in a modular manner to meet local or community requirements. We are implementing this protocol for arXiv.org and also provide an experimental prototype for the English Wikipedia as well as a client API.

6. ACKNOWLEDGMENTS

This work is supported through a generous grant from the Alfred P. Sloan Foundation, JISC, as well as a Marie Curie International Outgoing Fellowship (PIOF-GA-2009-252206) within the 7th European Community Framework Programme. We would like to thank Peter Kalchgruber for developing the ResourceSync Wikipedia prototype.

7. REFERENCES

- [1] B. Haslhofer and A. Isaac. data.europeana.eu - the europeana linked open data pilot. In *International Conference on Dublin Core and Metadata Applications - DC 2011*, 2011.
- [2] S. Hellmann, C. Stadler, J. Lehmann, and S. Auer. DBpedia live extraction. *On the Move to Meaningful Internet Systems: OTM 2009*, pages 1209–1223, 2009.
- [3] M. Klein, R. Sanderson, H. Van de Sompel, S. Warner, B. Haslhofer, C. Lagoze, and M. L. Nelson. A technical framework for resource synchronization. *D-Lib Magazine*, 19(1), 2013.
- [4] P. Mendes, M. Jakob, A. Garcia-Silva, and C. Bizer. DBpedia Spotlight: Shedding light on the web of documents. In *Proceedings of the 7th International Conference on Semantic Systems*, pages 1–8. ACM, 2011.
- [5] A. Tridgell and P. Mackerras. The rsync algorithm. Technical Report TR-CS-96-05, Australian National University, 1996. Available at: http://rsync.samba.org/tech_report/.
- [6] H. Van de Sompel, R. Sanderson, M. Klein, M. L. Nelson, B. Haslhofer, S. Warner, and C. Lagoze. A perspective on resource synchronization. *D-Lib Magazine*, 18(9), 2012.

⁹<http://arxiv.resourcesync.net>

¹⁰<http://en.wikipedia.resourcesync.net>