

Event Notifications in Value-Adding Networks

Patrick Hochstenbach^{1,2}[0000-0001-8390-6171], Herbert Van de Sompel^{2,3}[0000-0002-0715-6126], Miel Vander Sande^{4,5}[0000-0003-3744-0272], Ruben Dedecker²[0000-0002-3257-3394], and Ruben Verborgh²[0000-0002-8596-222X]

¹ Ghent University Library, Ghent University, Ghent, Belgium

² IDLab, ELIS, Ghent University - IMEC, Ghent, Belgium

³ DANS, Data Archiving and Networked Services, The Hague, The Netherlands

⁴ Meemoo, Vlaams Instituut voor het Archief, Ghent, Belgium

⁵ Dutch Digital Heritage Network, The Hague, The Netherlands

{Patrick.Hochstenbach,Herbert.VandeSompel,Ruben.Dedecker,Ruben.Verborgh}@UGent.be
Miel.VanderSande@meemoo.be

Abstract. Linkages between research outputs are crucial in the scholarly knowledge graph. They include online citations, but also links between versions that differ according to various dimensions and links to resources that were used to arrive at research results. In current scholarly communication systems this information is only made available post factum and is obtained via elaborate batch processing. In this paper we report on work aimed at making linkages available in real-time, in which an alternative, decentralised scholarly communication network is considered that consists of interacting data nodes that host artifacts and service nodes that add value to artifacts. The first result of this work, the “Event Notifications in Value-Adding Networks” specification, details interoperability requirements for the exchange of real-time life-cycle information pertaining to artifacts using Linked Data Notifications. In an experiment, we applied our specification to one particular use-case: distributing Scholix data-literature links to a network of Belgian institutional repositories by a national service node. The results of our experiment confirm the potential of our approach and provide a framework to create a network of interacting nodes implementing the core scholarly functions (registration, certification, awareness and archiving) in a decentralized and decoupled way.

Keywords: Scholarly communication · Digital libraries · Open science

1 Introduction

In the 2004 paper “Rethinking Scholarly Communication” [1], Van de Sompel et al introduced the perspective of an open scholarly communication system in which decentralized service hubs add value to research outputs hosted in repositories by fulfilling the core functions of scholarly communication - registration, certification, awareness, archiving [2] - in a decoupled manner and in ways they

see fit. The fact that value of a certain kind was added was intended to be openly exposed and, as such, would support creating bidirectional linkages between resources jointly involved in a life-cycle event, such as a preprint and a review thereof, an article and a copy of it in a long-term archive, etc. The authors understood that, in order for such an environment to be feasible, “the service hubs needed to be interconnected, as if they were part of a global scholarly communication workflow system”.

A similar perspective is at the basis of the more recent “Linked Research on the Decentralized Web” [3] and COAR’s “Next Generation Repositories report” [4] but remains in sharp contrast with the status quo. The scholarly communication system does not provide real-time life-cycle information. To a large extent, linkages between related research outputs are revealed by means of post-factum, heuristic-bound mining processes across research output corpora (e.g., Cabernac et al efforts to discover preprint/publication links [5]), that are so laborious that only the happy few can consider providing them as production services.

This paper introduces the “Event Notifications in Value-Adding Networks” [6] specification, which is a first step towards realizing the ability to expose life-cycle information for research outputs in real-time and hence to reveal linkages as they become known. In the network (see Figure 1) considered by the specification, service nodes add value to artifacts hosted by data nodes and the nodes exchange messages with that regard: a data nodes asks a service node to add value to one of its artifacts, and the service node responds with information about the service result; a service node volunteers information about value it unilaterally added to an artifact hosted by a data node. For example, a researcher can request that a research output be peer-reviewed via a certification service or directly by selected peers; she can request that a search engine include the research output in its index; she can ask an archive to collect the output and store it for posterity; a service can inform the researcher that her research output was leveraged by another endeavor; a peer can inform the researcher about a new annotation made to the research output. In the scholarly realm, the nature of the added value crucially includes fulfilling the core functions of scholarly communication. But, while scholarly communication was the main driver for the specification, it is conceived in a generic way and allows for a wide range of added-value as well as applications in other domains such as cultural heritage.




Fig. 1. Data nodes and service node exchanging notification messages.

The messages between the autonomous systems that operate on the network are point-to-point, owing to its decentralized nature. The “Event Notifications in Value-Adding Networks” specification details interoperability regarding the messaging protocol and the communication payloads that are essential to allow the envisioned network to be functional.

Our specification effort is generally motivated by the Decentralized Web [7] movement that emerged in reaction to the disconcerting evolution of the Web towards an environment in which a limited number of portals have become dominant, control personal data, and monetize online presence information as a product. Similar concerns exist in scholarly communication (e.g. [8–10]), with consolidation of services that cover the entire research life-cycle and publishers turning into data analytics companies. Decoupling content hosting from service provision yields a separation between ownership, control, and storage of data on the one hand, and the functionality that can be provided on top of it, on the other. It levels the playing field to allow for new entrants and creativity.

The remainder of this paper is structured as follows. Section 2 discusses related work. Our ongoing standards-based specification work is discussed in Section 3. Section 4 reports on an experiment in which data-literature links are pushed towards the repositories that host the resources involved in such links in a manner that complies with the specification. Section 5 provides a summary and an outlook for future work.

2 Related Work

With `dokie.li` [11], Capadisli introduced the use of Linked Data Notifications [12] (LDN) with `ActivityStreams2` [13] (AS2) payloads as a push-oriented interoperability glue for communication among researchers that publish research outputs in personal data pods and with services that add value to these outputs. His approach provided inspiration for the Researcher Pod [14] and COAR Notify [15] projects.

The Researcher Pod project, funded by the Andrew W. Mellon Foundation, explores a scholarly communication system based on the Decentralized Web vision of Tim Berners-Lee [16] as was envisioned by Capadisli [3] and Van de Sompel [17]. A Researcher Pod network is highly decentralized: each researcher publishes research outputs in a personal data pod [18, 19] that is hosted in a personal web domain. Value is added to these research outputs through interaction with appropriate decoupled services and/or peers. The Researcher Pod network uses LDN and AS2 as core interoperability ingredients, but the AS2 payloads are strictly profiled in order to avoid ambiguous interpretation of notifications. Our “Event Notifications in Value-Adding Networks” specification is one of the outcomes of this project.

In COAR Notify, the network consists of repositories and providers of overlay review services. A core goal of the project is establishing bidirectional links between preprints and certifications thereof (e.g., stand-alone reviews, endorsements, peer-reviewed articles) as soon as the latter get published. Our specifica-

tion work is closely aligned with COAR Notify to which we actively contribute, uses the same interoperability ingredients, but has a broader scope regarding what constitutes added value and in which communities the specification can be used.

ScholarXplorer [20], based on the Scholix framework [21, 22], is a centralized approach to exchange life-cycle information about research outputs, focused on the discovery of research data-literature links. Scholix was partly motivated by the lack of interoperability among various approaches for the point-to-point exchange of data-literature links that had resulted from numerous bilateral agreements [23]. It introduced an approach in which a limited number of hubs, including Crossref, DataCite, and OpenAIRE collect and mine linkage information from a variety of scholarly and web corpora. The hubs support retrieving the link information, which is uniformly expressed according to a link information package schema [24]; no specific protocol for transporting these packages is selected. ScholarXplorer collects the link information from the hubs and makes it available by means of a query-oriented API allowing, for example, a repository to recurrently poll ScholarXplorer in search of linkages in which its resources are involved. Our specification details an approach that allows expressing and exchanging data-literature (and other) links in a push-oriented, point-to-point manner that leverages W3C standards. Conceptually, it aims to provide the capability to broadcast linkage information directly between the systems that host the artifacts as soon as a link is established. This in contrast with the Scholix framework that needs to wait for scholarly artifacts to be published, harvested and released in batches, months after the artifact creation date.⁶

3 The Value-Adding Network

3.1 Design and Technology Considerations

The Value-Adding Network approach utilizes a *point-to-point* communication style, avoiding the need for centralized hubs that can eventually become too powerful to remain trustworthy and too big to fail. The approach is *push-oriented*, with only the relevant nodes being updated about new information as soon as it becomes available. In the highly distributed Value-Adding Network set-up, a pull-oriented approach would introduce a hardly tractable “which node to pull” challenge in addition to the well-known “when to poll” conundrum. Interactions are necessarily *asynchronous* owing to the fact that the time between requesting a service and providing a service is unpredictable and could range between almost immediate (e.g., provide a trusted timestamp) to months (e.g., provide a peer-review report).

The communication patterns in the Value-Adding network are *action-oriented*. The patterns express *when* an activity was initiated, acknowledged, or yielded a result (See Section 3.3). The internal workflow of *how* the activity was executed

⁶ See for example <https://scholarxplorer.openaire.eu/#/statistics> for an overview of ScholarXplorer batch frequencies

is treated as a black box. As such, message payloads contain only core information (in most cases URL identifiers) to convey in which value-adding activities data nodes, service nodes, and artifacts are involved and what the service results are. The assumption is that further information about the entities involved is available by using their identifiers in auto-discovery mechanisms supported by protocols such as Signposting [25], WebID [26], Solid [27], and Hydra [28]. As such, message payloads can be expressed in a *by reference* rather than *by value* style.

Several candidate protocols exist that could meet the point-to-point, push-oriented, and asynchronous design choices, including WebHooks [29], Webmention [30], and Linked Data Notifications [12]. Webhooks, while commonly used, is not formally standardized and does not have a natural tie-in with specific serializations and/or vocabularies. Webmention is uni-directional and, as such, can't be used off-the-shelf for a request-response communication style. Moreover, it does not actually allow transmitting messages and rather merely informs the target resource that the sender resource mentions (links to) it. We chose for Linked Data Notifications in combination with the ActivityStreams 2.0 vocabulary, both W3C standards. That combination was successfully used in the dokie.li work, and has proven to work at scale with ActivityPub [31] implementations that provide a communication channel for millions of users on social networks such as Mastodon [32].

3.2 A network of Data Nodes and Service Nodes

In a scholarly Value-Adding Network, artifacts are research outputs hosted by **Data Nodes**, while **Service Nodes** are systems that add value to these research outputs. Both types of node are essentially web servers that host one or more LDN Inboxes where **Notifications** can be received. An LDN Inbox is an HTTP endpoint associated with a web server to which notifications can be POSTed. Our specification describes how a value-added service can be triggered by sending an appropriate notification to LDN Inbox of a service node. And it describes how the outcome of providing a value-adding service for an artifact, the **Service Result**, can be communicated to a data node's LDN Inbox by means of a notification. The notification patterns and payload are discussed in Section 3.3 and Section 3.4, respectively.

The manner in which value-added services are provided is not in scope of our specification. The internal workflows of a service node is considered as a black box yet could be openly exposed on the basis of other specifications. How nodes, or better the agents that operate on behalf of the nodes, act upon receipt of a notification is also out of scope of our specification. For instance, the notifications could be added to a queue for further processing, e.g., to update the metadata of an artifact, or to write life-cycle information in an **Event Log** [33] that is made accessible to support the creation of scholarly knowledge graphs that cover artifacts across data nodes.

3.3 Network communication patterns

At this point, our Value-Adding Network specification details two communication patterns: a one-way pattern and a request-response pattern.

One-way pattern. This pattern is used when the agent that operates on behalf of a node sends an unsolicited message about an artifact. One-way patterns are used for sending informational messages pertaining to an artifact to another node in the network. These messages can originate from a data node or a service node. Examples of one-way patterns are:

- A repository (data node) informs a subject index (service node) about a new artifact.
- A linkage service (service node) informs a repository (data node) about a discovered linkage between a hosted data set artifact and a publication.

Request-response pattern. This pattern allows agents that act on behalf of a data node to ask a service node to apply a value-added service to one of its artifacts. There is no expectation of an immediate response but definitely the expectation of an eventual response. The pattern doesn't assume there is only one response (or request) but introduces a communication thread (similar to email threads) which will be used to relate notifications to each other. An example of a request-response pattern is:

- A repository (data node) would like a peer review for one of the artifacts and therefore contacts a peer review service (service node) to provide this service. The peer review service can first evaluate the request and acknowledge or reject it. When acknowledged, a peer review can be made available at a later date, and a reference to this peer review can be sent back in a notification to the repository.

3.4 Anatomy of a notification message

Notifications are exchanged in order to inform agents about added-value that is requested or was provided for artifacts. Network agents send and receive the notifications, data and service nodes provide LDN Inboxes to accept incoming messages. Value-Adding Networks uses a subset of the ActivityStreams 2.0 vocabulary [13] serialized as JSON-LD [34] payloads. The LDN protocol requires all the payloads to be expressed by default as JSON-LD, but other RDF serializations are allowed. For brevity, all examples in this paper use Turtle syntax [35]. We refer to our specification for further JSON-LD examples.

Every notification has at least one AS2 type that is a subset of the `as:Activity` types. The subtypes are limited to those that express facts about artifacts and currently exclude types related to typical social activities (such as `as:Like`, `as:Follow`). The subtypes used are:

- Activities to request or receive notifications about value adding services: `as:Announce`, `as:Offer`, `as:Accept`, `as:Reject`, and `as:Undo`
- Activities that express CRUD (*Create, Update, Delete*) life-cycle event on data nodes: `as:Create`, `as:Update`, and `as:Remove`.

In addition to the activity types listed above, an activity may have additional subtypes that are specific to an application domain. For instance, when a service node would like to use `as:Announce` for the endorsement of a publication an extra `schema:EndorseAction` could be added as activity type.

Table 1 shows the mapping between AS2 entities and the Value-Adding Network entities that are used for the notification shown in Listing 1, which illustrates a typical notification that uses the **one-way pattern** from a service node to a data node. In it, an agent at the end of the service node uses the `as:Announce` activity type to inform an agent at the end of a data node about the result of adding value to one of the latter’s hosted artifacts. In human language, the notification expresses the following: “Fairfield Archive (service node) announces to Springfield Library (data node) a service result concerning Springfield’s artifact <https://springfield.library.net/artifact/13-02.html>. For this artifact, the <https://fairfield.org/archive/version/317831-13210> is a memento.”

Table 1. One-way AS2 pattern for notification messages from a service node to a data node.

AS2 element	Description
<code>id</code>	Message identifier
<code>type</code>	Activity type
<code>as:actor</code>	Agent that performed the activity
<code>as:origin</code>	Agent responsible for sending the notification
<code>as:context</code>	The artifact on the data node for which an value-added service was provided
<code>as:object</code>	The result of the value-added service provided for an artifact on the data node
<code>as:target</code>	The agent at the data node that is the addressee of the notification

```

@prefix as: <https://www.w3.org/ns/activitystreams#> .
@prefix ldp: <http://www.w3.org/ns/ldp#> .

<urn:uuid:239FD510-03F4-4B56-B3A0-0D3B92F3826D> a as:Announce ;
  as:actor <https://fairfield.org/about#us> ;
  as:origin <https://fairfield.org/system> ;
  as:context <https://springfield.library.net/artifact/13-02.html> ;
  as:object <urn:uuid:CF21A499-1BDD-4B59-984A-FC94CF6FBA86> ;
  as:target <https://springfield.library.net/about#us> .

<https://fairfield.org/about#us> a as:Organization ;
  ldp:inbox <https://fairfield.org/inbox> ;
  as:name "Fairfield Archive" .

<https://fairfield.org/system> a as:Service;
  as:name "Fairfield Archive System".

<urn:uuid:CF21A499-1BDD-4B59-984A-FC94CF6FBA86> a as:Relationship ;
  as:object <https://fairfield.org/archive/version/317831-13210> ;
  as:relationship <https://www.iana.org/memento> ;
  as:subject <https://springfield.library.net/artifact/13-02.html> .

```

```
<https://springfield.library.net/about#us> a as:Organization;
  ldap:inbox <https://springfield.library.net/inbox/> ;
  as:name "Springfield Library" .
```

Listing 1. A value-adding Activity Streams 2.0 notification using Turtle syntax

4 Experimental Investigation

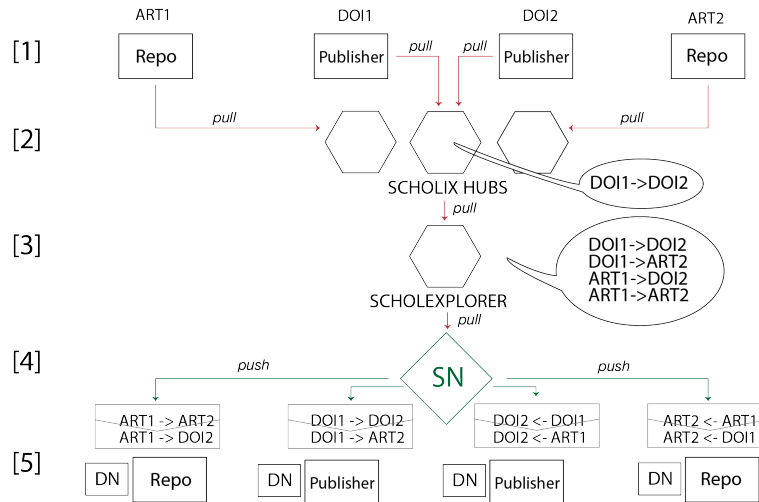


Fig. 2. Scholix network with an added national service node.

The COAR Notify project has already illustrated the feasibility of real-world use of profiled LDN/AS2 notifications in a request-response pattern between repositories and certification services. We set up an experiment to investigate the one-way pattern where a service node informs a network of data nodes about the *linkages* their artifacts are involved in. We leverage the existing Scholix link base which provides an extensive collection of such linkage information.

In the Scholix network, data-literature linkages are provided via a limited number of hubs such as DataCite, CrossRef and OpenAIRE ([2] in Fig. 2), that obtain information from systems that host research outputs ([1] in Fig. 2). ScholeXplorer ([3] in Fig. 2) is an aggregation service that collects link information from the hubs and makes the result linkages available as a dataset at an irregular frequency. ScholeXplorer can be polled to discover data-literature links between research outputs hosted in the systems covered by the hubs. For our experiment, we prototype a national service NatServ ([4] in Fig. 2) that queries ScholeXplorer for data-literature links in which artifacts hosted by Belgian institutional repositories are involved, and that pushes the link information

to the relevant repository according to the Value-Adding Network approach. Using the terminology of that specification, NatServ is a service node that pushes link information to data nodes associated with the institutional repositories ([5] in 2). Figure 2 also shows the research outputs identified by ART1, DOI1, DOI2, ART2 hosted by a variety of repositories. The hubs pull information about these artifacts, and as a result of data processing, determine that DOI1 is related to DOI2. ScholeXplorer pulls such linkage information from all hubs and further discovers that DOI1 is a manifestation of ART1 and DOI2 is a manifestation of ART2. Finally, our prototype NatServ (SN) pulls data-literature links pertaining to selected repositories from ScholeXplorer and pushes the information using profiled LDN/AS2 notifications to the LDN Inboxes of data nodes associated with the pertinent repositories, e.g. the data node at the left bottom receives information that its ART1 is related to ART2 and DOI2.

4.1 Experiment

We downloaded Scholix data for the following Belgian repositories: *Institutional Repository Universiteit Antwerpen* (Antwerpen) , *Ghent University Academic Bibliography* (Biblio) and *Open Repository and Bibliography - University of Liège* (Orbi). These downloads resulted in 711 records for Antwerpen, 1056 for Biblio, and 669 for Orbi. Between May 10 and May 25 2022, we also tried to download Scholix data for the KU Leuven *Lirias* repository and the Vrije Universiteit Brussel *Vrije Universiteit Brussel Research Portal* repository. Although hundreds of Scholix records are available that indicate linkages in which their research outputs are involved, none contained actual URLs of those artifacts.⁷ Since these URLs are needed for the remainder of the experiment, we excluded the Scholix records for these repositories. At the time of writing, our communication with ScholeXplorer representatives have not led to an explanation for this problem.

Discovery of LDN Inboxes With NatServ in possession of the linkages in which artifacts of the three repositories are involved, its next step is to discover the data nodes and associated LDN Inboxes to which linkage information needs to be pushed. The goal is to push the information to the systems that host source as well as target of the links. Since we assume auto-discovery mechanisms for the discovery of pertinent Inboxes, this requires dereferencing the URLs of the all artifacts (expressed in Scholix as PID-URLs, e.g. DOI⁸, HANDLE⁹, PMID¹⁰, PMC¹¹, and ARXIV¹²) up to their landing pages and extracting the link with <http://www.w3.org/ns/ldp#inbox> Link relation from its HTTP header.¹³ The

⁷ see e.g. <https://gist.github.com/phochste/4d6ed7f7748c5d74417ac45d38480540>

⁸ Digital Object Identifier <https://doi.org>

⁹ Handle identifier <http://hdl.handle.net>

¹⁰ PubMed Identifier <https://pubmed.ncbi.nlm.nih.gov>

¹¹ PubMed Central <https://www.ncbi.nlm.nih.gov/pmc/>

¹² arXiv Identifier <https://arxiv.org>

¹³ See <https://www.w3.org/TR/ldn/#discovery>

URL dereferencing process is less trivial than it may seem [36] and we employed the simplest possible algorithm by following for each link identifier all the HTTP redirects until a page with a HTTP 200 response was found. Obviously, no LDN Inboxes were found because none of the systems involved in linkages currently support Value-Adding Networks, but the exercise remains important as a means to understand the effort that would be involved if they did.

Table 2 provides an overview of the complete resolved data-literature link network for each Belgian institution: this is the network of all data-literature links involving artifacts hosted by the institutional repositories. For instance, if one institutional artifact references ten datasets, then this is a network of 10+1 artifact URLs.

Table 2. Number of artifact URLs resolved for the data-literature network of each Belgian institution and time required to resolve PID-URLs to their landing page.

Scholix Link Provider	#Records	# Artifact URLs	#Resolve time (sec)	time/req
Antwerpen	711	4335	695	0.978 ± 0.01 s
Biblio	1056	7189	3651	3.457 ± 0.02 s
Orbi	669	3375	367	0.549 ± 0.02 s

To proceed with the experiment, we created LDN Inboxes for all systems that host artifacts involved in the linkages, using a Solid CSS server as proxy of LDN Inboxes.¹⁴ We created the inboxes at addresses generated as follows `https://{data.host}/inbox` where `data.host` is the base url of the landing page. For example, the LDN Inbox URL: `https://arxiv.org/inbox` will be hosted at: `http://solidpod-baseurl/arxiv.org/inbox` where `solidpod-baseurl` is the base URL of the Solid CSS server.

With Solid CSS, the implementation of LDN Inboxes doesn't require any fundamental changes in institutional repositories to add read-write capabilities to an existing network. In our experiment effortlessly created hundreds of LDN Inboxes given the fact that: 1) LDN is a subset of the Solid protocol stack [18], and, 2) an LDN Inbox can be put on any node as long as it can be discovered from the artifact landing page.

In total, 197 [Antwerpen], 245 [Biblio], and 193 [Orbi] LDN Inboxes were created to cover the data-literature network of the institutions.

Generating notification messages Next, we use the guidelines from our specification for sending data about a linking event to transform a ScholeXplorer data-literature link into a profiled LDN/AS2 payload.¹⁵ The generated payloads are nimble (see snippet in Listing 2) because we assume that further information pertaining to entities involved can be obtained via auto-discovery mechanisms.

¹⁴ <https://github.com/CommunitySolidServer/CommunitySolidServer>

¹⁵ <https://www.eventnotifications.net/#pattern-5-a-service-node-notifies-a-data-node-about-linking-event>

The ActivityStreams 2.0 Activity `as:Relationship` is used to express a relationship between two artifacts and the mapping that was used is shown in Table 3.¹⁶

Table 3. One-way AS2 mappings for notification messages in the Scholix experiment.

AS2 element	Description
<code>as:actor</code>	We minted the WebID <code>https://scholixexplorer.openaire.eu/#about</code> .
<code>as:origin</code>	We minted the WebID <code>https://mellonscholarlycommunication.github.io/about#us</code>
<code>as:context</code>	Points to the artifact URL on the target data node the data-literature link is about.
<code>as:object</code>	Contains the description data-literature link. An example such a data-literature link description is shown in Listing 2.
<code>as:target</code>	This contains the generated (or discovered) LDN Inbox for each data-literature link.

```

1 <urn:uuid:240c0091-b271-4e44-87f7-5598da5b24ad>
2   a as:Relationship ;
3   as:object <https://dx.doi.org/10.5061/dryad.10hq7> ;
4   as:relationship <http://www.scholix.org/References> ;
5   as:subject <https://biblio.ugent.be/publication/8646849> .

```

Listing 2. Snippet of the `as:object` of data-literature notification.

Sending notifications to LDN Inboxes For each institution, we sent the generated notifications to the simulated LDN Inboxes of all systems for which artifacts are involved in links either as source and target. This doubles the number of messages that is provided in Table 2. For instance, if one artifact *A* references two data sets *B* and *C*, then we send two messages for *A*: *A* references *B*, *A* references *C*. To *B* we can send: *A* references *B*. To *C* we can send *A* references *C*. In total four messages.

Table 4 shows the number of messages that were posted and the typical throughput that is not dependent on the payload of the message.

The generated inboxes and the notifications that were sent to them can be downloaded at Zenodo.¹⁷ The software to download and transform Scholix into AS2 notifications and to publish notifications to LDN Inboxes is available at GitHub.¹⁸ All experiments were run on a 2 CPU, 1 GB, Linux host connected to the Belnet research network.

¹⁶ The full Turtle representation of the notification message is available online <https://gist.github.com/phochste/afdc7bc5d6a09cec2b12c60755ef1c82>

¹⁷ <https://zenodo.org/record/6555821>

¹⁸ <https://github.com/MellonScholarlyCommunication/scholix-client>

Table 4. Sending LDN Notifications for the complete network of three Belgian institutions. The mean posting time for these networks have a constant rate of about 80 notifications per second.

Scholix Link Provider	# Sent Notifications	#Post time (sec) & time/req
Antwerpen	8670	108s , 80 req/sec
Biblio	14378	183s , 78 req/sec
Orbi	6720	86s , 78 req/sec

5 Summary and Outlook

Our work demonstrated how the introduction of decentralized read-write technology on top of existing research networks provides a push-based mechanism to provide up-to-date information of life-cycle events pertaining to artifacts that are stored in the network. In our case, the network was applied to the distribution of data-literature links by an experimental national NatSev service node (built on top of ScholeXplorer) to data nodes in the network. These experiments demonstrate how the service node was able to inform a network of institutional repositories about data-literature links using push-based LDN Notifications.

The scalability of the distribution of data-literature links in the Scholix experimental network is dependent on the time it takes to resolve PID-URLs to the landing pages of the artifact. Even with our naive resolve algorithm, which doesn't employ any optimization strategies, the complete Belgian national data-literature network for all known linkages could be distributed within two hours on a small Linux host.

Our Scholix experiment was dependent on data mining capabilities of centralized Scholix hubs. The real goal of a distributed Value-Added Network is for linkages to be mined (or cataloged), and sent from the repository that hosts an artifact that is the source of a link to a repository that hosts the artifact that is the target of that link. Research in our Researcher Pod project has already been started to explore how data-mining capabilities can be added to Value-Added networks via orchestration [39].

With our small sample of Scholix providers, we could already create for three Belgian institutions a network with $197 + 245 + 193$ data nodes. It would be a bright future for the decentralisation of scholarly communication when all linkage information could be made available in real-time (instead of months after a long publication and review process) and point-to-point instead of requiring a few massive central hubs and gateways. The question how this can be accomplished in a repository landscape that has hardly invested in interoperability affordances since the broad adoption of OAI-PMH [40] in the early 2000s.

Our hope is that COAR's "Next Generation Repositories" vision will get a broader uptake. Important steps are taken in the COAR Notify project with the first concrete implementations that make profiled LDN/AS2 notifications a reality for communication between repositories and review services. Arcadia has recently recommended a 4 year, US\$4 million grant [41] to accelerate COAR Notify enhancing the role of repositories, thereby transforming scholarly commu-

nications, making it more research-centric, community-governed, and responsive to the diverse needs across the globe.

Acknowledgements This work is funded by the Andrew W. Mellon Foundation (grant number: 1903-06675) and supported by SolidLab Vlaanderen (Flemish Government, EWI and RRF project VV023/10). The authors would like to thank the COAR Notify (Kathleen Shearer, Paul Walk, Martin Klein) for involving them from the outset, allowing them to provide input and gain valuable insights for their generic specification effort.

References

1. Van de Sompel, H., Payette, S., Erickson, J., Lagoze, C., Warner, S. : Rethinking Scholarly Communication. *D-Lib Magazine*, 10(9) (2004) <https://www.dlib.org/dlib/september04/vandesompel/09vandesompel.html>
2. Roosendaal, H., Geurt, P.: Forces and functions in scientific communication: an analysis of their interplay. (1998) <https://perma.cc/5HYM-BEKF>. Captured 3 June 2019
3. Capadisli, S.: Linked Research on the Decentralised Web. (2020) <https://csarven.ca/linked-research-decentralised-web>
4. Rodrigues, E., Bollini, A., Cabezas, A., Castelli, D., Carr, L., Chan, L., Humphrey, C., Johnson, R., Knoth, P., Manghi, P., Matizirofa, L., Perakakis, P., Schirrwagen, J., Selematsela, D., Shearer, K., Walk, P., Wilcox, D., Yamaji, K.: Next Generation Repositories: Behaviours and Technical Recommendations of the COAR Next Generation Repositories Working Group. Zenodo. (2017) <https://doi.org/10.5281/zenodo.1215014>
5. Cabanac, G., Oikonomidi, T. & Boutron, I. Day-to-day discovery of preprint-publication links. *Scientometrics* 126, 5285–5304 (2021). <https://doi.org/10.1007/s11192-021-03900-7>
6. Hochstenbach, P., Vander Sande, M., Dedecker, R., Walk, P., Klein, M., Van de Sompel, H. : Event Notifications in Value-Adding Networks. (2020). <https://www.eventnotifications.net> Accessed 14 Jul 2022
7. Berners-Lee, T : Socially Aware Cloud Storage. (2009) <https://www.w3.org/DesignIssues/CloudStorage.html>
8. Pooley, J. : Surveillance Publishing. Elephant in the Lab. (2022) <https://doi.org/10.5281/zenodo.6384605>
9. Siems, R. : Das Lesen der Anderen: Die Auswirkungen von User Tracking auf Bibliotheken. *O-Bib. Das Offene Bibliotheksjournal / Herausgeber VDB*, 9(1), 1–25. (2022) <https://doi.org/10.5282/o-bib/5797>
10. Posada, A., Chen, G.: Inequality in Knowledge Production: The Integration of Academic Infrastructure by Big Publishers. *ELPUB 2018*, Jun 2018, Toronto, Canada. [dx.doi.org/10.4000/proceedings.elpub.2018.30](https://doi.org/10.4000/proceedings.elpub.2018.30)
11. Capadisli, S.: Dokieli <https://dokie.li>
12. Capadisli, S., Guy, A.: Linked Data Notifications. W3C Recommendation. (2017) <https://www.w3.org/TR/ldn/>
13. Snell, J., Prodromou, E.: Activity Streams 2.0. W3C Recommendation. (2017) <https://www.w3.org/TR/activitystreams-core/>

14. "Scholarly Communications in the Decentralized Web". <https://mellon.org/grants/grants-database/grants/ghent-university/1903-06675/>. Accessed 20 May 2022
15. "The Notify Project". <https://www.coar-repositories.org/notify/>. Accessed 20 May 2022
16. Berners-Lee, T.: <https://www.w3.org/DesignIssues/Principles.html>.
17. Van de Sompel, H.: Scholarly Communication: Deconstruct & Decentralize? <https://www.youtube.com/watch?v=o4nUe-6Ln-8>
18. Sambra, A.V., Mansour, E., Hawke, S., Zereba, M., Greco, N., Ghanem, A., Zagidulin, D., Abounaga, A., Berners-Lee, T.: Solid: a platform for decentralized social applications based on linked data. Tech. rep., Technical report, MIT CSAIL & Qatar Computing Research Institute (2016) http://emansour.com/research/lusail/solid_protocols.pdf
19. Mansour, E., Sambra, A.V., Hawke, S., Zereba, M., Capadisli, S., Ghanem, A., Abounaga, A., Berners-Lee, T.: A demonstration of the solid platform for social web applications. In: Proceedings of the 25th International Conference Companion on World Wide Web. pp. 223–226 (2016) <https://dl.acm.org/doi/10.1145/2872518.2890529>
20. La Bruzzo, S., Manghi, P. : OpenAIRE ScholeXplorer Service: Scholix JSON Dump (5.0) [Data set]. Zenodo. (2022) <https://doi.org/10.5281/zenodo.6338616>
21. Burton, A., Koers, H., Manghi, P., Stocker, M., Fenner, M., Aryani, A., La Bruzzo, S., Diepenbroek, M. and Schindler, U. : The Scholix framework for interoperability in data-literature information exchange. D-Lib Magazine, 23(1/2) (2017) <http://www.dlib.org/dlib/january17/burton/01burton.html>
22. Burton, A., Koers, H., Manghi, P., La Bruzzo, S., Aryani, A., Diepenbroek, M., Schindler, U. : On Bridging Data Centers and Publishers: The Data-Literature Interlinking Service. In: Garoufallou, E., Hartley, R., Gaitanou, P. (eds) Metadata and Semantics Research. MTSR 2015. Communications in Computer and Information Science, vol 544. Springer, Cham. (2015) <https://dx.doi.org/10.1108/PROG-06-2016-0048>
23. Callaghan, S., Tedds, J., Lawrence, R., Murphy, F., Roberts, T., and Wilcox, W. : Cross-Linking Between Journal Publications and Data Repositories: A Selection of Examples. International Journal of Digital Curation. (2014) <https://doi.org/10.2218/ijdc.v9i1.310>
24. "Scholix metadata schema for exchange of scholarly links". <http://www.scholix.org/schema> Accessed 24 May 2022
25. "Signposting the Scholarly Web". <https://signposting.org>. Accessed at 20 May 2022
26. W3C. "WebID". <https://www.w3.org/wiki/WebID>. Accessed at 25 May 2022
27. Capadisli, S., Berners-Lee, T., Verborgh, R., Kjensmo, K. : Solid Protocol Version 0.9.0, 2021-12-17. <https://solidproject.org/TR/protocol> Accessed at 25 May 2022
28. Lanthaler, M.: Hydra Core Vocabulary. W3C Unofficial Draft (2021) <https://www.hydra-cg.com/spec/latest/core/>
29. "Webhooks – The Definitive Guide". (2022) <https://webhook.net>. Accessed at 20 May 2022
30. Parecki, A. : Webmention . W3C Recommendation. (2017) <https://www.w3.org/TR/webmention/>
31. Lemmer-Webber, C., Tallon, J., Shepherd, E., Guy, A., Prodromou, E.: Activity-Pub. W3C Recommendation. (2019) <https://www.w3.org/TR/activitypub/>
32. "Giving social networking back to you - Mastodon". <https://joinmastodon.org>. Accessed at 20 May 2022

33. Vander Sande, M., Hochstenbach, P., Dedecker, R., Werbrouck, J. : *Artefact Lifecycle Event Log*. (2022) <https://mellonscholarlycommunication.github.io/spec-eventlog/>
34. Sporny, M., Longley, D., Kellogg, G., Lanthaler, M., Champin, P.A., Lindström, N. : *JSON-LD 1.0. W3C Recommendation*. (2014) <https://www.w3.org/TR/json-ld/>
35. Beckett, D., Berners-Lee, T., Prud'hommeaux, E., Carothers, G. : *RDF 1.1 Turtle* . W3C Recommendation. (2014) <https://www.w3.org/TR/turtle/>
36. Klein, M., Balakireva, L.: *On the Persistence of Persistent Identifiers of the Scholarly Web*. In: Hall, M., Merčun, T., Risse, T., Duchateau, F. (eds) *Digital Libraries for Open Knowledge. TPD L 2020. Lecture Notes in Computer Science()*, vol 12246. Springer, Cham. (2020) https://doi.org/10.1007/978-3-030-54956-5_8
37. Hochstenbach, Patrick : *Notifications in Value-Adding Networks (1.0.0)*. Zenodo. (2022). <https://doi.org/10.5281/zenodo.6555821>
38. Hochstenbach, P.: *Mellon Scholarly Communication*, (2022), GitHub repository. <https://github.com/MellonScholarlyCommunication/scholix-client>
39. Vander Sande, M., Hochstenbach, P., Dedecker, R., Werbrouck, J.: *Orchestrator for a decentralized Web network. Working Draft*, 26 November 2021 (2021) <https://mellonscholarlycommunication.github.io/spec-orchestrator/>
40. Lagoze, C., Van de Sompel H., Nelson, M., Warner, S. : *The Open Archives Initiative Protocol for Metadata Harvesting. Protocol Version 2.0 of 2002-06-14* (2015) <http://www.openarchives.org/OAI/openarchivesprotocol.html>
41. COAR. "COAR welcomes significant funding for the Notify Project". <https://www.coar-repositories.org/news-updates/coar-welcomes-significant-funding-for-the-notify-project/>. Accessed 25 May 2022