

The International Journal of Digital Curation

Issue 2, Volume 2 | 2007

The UK LOCKSS Pilot Programme: A Perspective from the LOCKSS Technical Support Service

Adam Rusbridge,

Digital Curation Centre & Humanities Advanced Technology and Information Institute
(HATII), University of Glasgow

Seamus Ross,

Digital Curation Centre & Humanities Advanced Technology and Information Institute
(HATII), University of Glasgow

December 2007

Abstract

Over the last decade libraries have increasingly shifted journal access from print to digital. The preference of users for online content, the demand of readers for a broader range of content, and the rising costs of library shelf space all contributed to bringing about this change. A variety of approaches has emerged to support access to these digital journals. The common ones require libraries to access content through publishers, which invariably means libraries lose long-term control of the journals to which they have subscribed. As an alternative, the Lots of Copies Keep Stuff Safe (LOCKSS) approach enables libraries to regain custody of journal assets while maintaining the access and licence terms stipulated by the publisher. This article describes the UK LOCKSS Pilot Programme; a two-year JISC/CURL (Joint Information Systems Committee/Consortium of Research Libraries in the British Isles)-funded initiative to explore issues associated with the practical implementation of LOCKSS in the UK Higher Education institutions. It examines the pilot from the perspective of the LOCKSS Technical Support Service, a core component of the Pilot Programme.



Introduction

Over the last decade libraries have increasingly shifted journal access from print to digital. The preference of users for online content, the demand of readers for a broader range of content, and the rising costs of library shelf space all contributed to bringing about this change. However, current publisher distribution models require libraries to subscribe to journal content accessed on a centralised publisher-maintained server; a model whereby libraries lease rather than own content. The Lots of Copies Keep Stuff Safe (LOCKSS¹) approach enables libraries to regain custody of journal assets while maintaining the access and licence terms stipulated by the publisher. This article describes the UK LOCKSS Pilot Programme²; a two-year JISC/CURL (Joint Information Systems Committee/Consortium of Research Libraries in the British Isles)-funded initiative to explore issues associated with the practical implementation of LOCKSS in the UK Higher Education institutions. This article examines the pilot from the perspective of the LOCKSS Technical Support Service, a core component of the Pilot Programme. After introducing the rationale for LOCKSS and the establishment of the UK LOCKSS Pilot Programme, we will describe some of the issues that have emerged and suggest ways in which we expect to move forward.

Changing Models of Electronic Journal Access

Web publishing provided ease of access, faster dissemination, and cost and space benefits to libraries. The development of this distribution and access model resulted in two significant by-products. Firstly, publishers wished to limit access to users who had active subscriptions or who were members of a subscribing institution. They were able to achieve those objectives by requiring user authentication to a central server. Secondly, there are many licensing options under which content may be accessed, such as annual subscriptions, bulk basket deals, short-term back-file access and aggregators. Different terms and conditions often apply to each licence type. Dynamically updating the access and usage conditions was reasonably easily achieved with a single, centralised service.

Journal users and librarians were concerned that this centralised model was characterised by a single point of failure. It lacked the redundancy inherent in the physical model, with copies distributed across many libraries which provided some measure of security against loss of access to journal content. In the digital environment what happens if a publisher ceases operation or a journal changes hands? How can libraries minimise the risks associated with the complexities, fragility, and interdependencies associated with digital materials? Librarians, moreover, expressed concern at the implications of leasing rather than owning content, and were keen to see mechanisms put in place to assure continued access by libraries to content for which they had paid (Waters et al., [2005](#)).

To address this, various journal archiving approaches have been developed and trialled (Kenney et al., [2006](#); Jones, [2007](#)). Each of these approaches has advantages and disadvantages. Legal deposit is limited to material relevant to the applicable country and access to legal deposit depots may be heavily restricted, often either to on-

1 LOCKSS Website, <http://www.lockss.org>

2 UK LOCKSS Pilot Strand, http://www.jisc.ac.uk/whatwedo/programmes/programme_preservation/programme_lockss.aspx

site access or a very small number of concurrent points of access (Boulderstone, [2007](#); Oltmans & van Wijngaarden, [2006](#)). Third-party non-profit archiving services (Fenton, [2006](#)) and community-driven “dark archive” approaches (Reich, [2006](#)) are emerging as possible options.

The LOCKSS system, being developed at Stanford University Libraries since 1999, is open source software that enables libraries to collect, maintain, and access local copies of web-published content. It supports the establishment of individual archives within each participating library, enabling the development of a persistent, well-managed collection of content relevant to the objectives of the participating libraries. As a result the libraries own rather than lease this content (Reich & Rosenthal, [2001](#)). A detailed discussion on the architecture of the LOCKSS system can be found in Rosenthal ([2003](#)), and the audit mechanism used for data monitoring and repair in Maniatis et al. ([2003](#)).

The UK LOCKSS Pilot Programme

The introduction of archiving clauses in the UK National Initiative for the Licensing of Electronic Journals (NESLi2) Model License³ was a first step in addressing the concerns of librarians about the loss of access to content and of publishers to the loss of control. Having outlined model terms for archiving journal content, efforts turned towards the exploration of practical methods through which the spirit of these model terms could be implemented. The LOCKSS approach offered an attractive option because it provides libraries with a cost-effective model that made it possible for them to develop a local, durable archive of electronic journal content. While LOCKSS is a distributed technical architecture, its implementation within the UK academic library environment required an element of centrally coordinated effort. JISC and CURL defined the three core aspects of the initiative that needed to be in place:

- a technical support service to manage the provision of support and coordination with the LOCKSS team in the United States,
- a negotiation agent to secure publishers’ agreement on the aim of LOCKSS to support the archiving of their content, and,
- a network of participating institutions.

In late 2005, JISC issued a call in partnership with CURL inviting libraries at UK higher education institutions to participate in a UK LOCKSS Pilot Programme⁴. The two-year pilot was launched in late February 2006 with 24 libraries participating initially⁵. In response to the high level of demand from the community to participate in the pilot, an additional six institutions were invited to join the pilot as Associate Members from July 2006. The Digital Curation Centre⁶ (DCC) through its partner the Humanities Advanced Technology and Information Institute⁷ (HATII) at the

3 NESLi2 is the UK’s national initiative for the licensing of electronic journals on behalf of the higher and further education and research communities, 2003-2006. See <http://www.nesli2.ac.uk> and http://www.nesli2.ac.uk/NESLi2_Licence.doc for the full text of the Model Licence, paying note to clauses 2.2.2, 5.4.1 and 5.4.2

4 Funding Circular 7/05: UK LOCKSS Pilot Programme http://www.jisc.ac.uk/fundingopportunities/funding_calls/2005/12/funding_circular7_05.aspx

5 For the full list of participating institutions, see: http://www.jisc.ac.uk/whatwedo/programmes/programme_preservation/programme_lockss.aspx

6 Digital Curation Centre, <http://www.dcc.ac.uk>

7 Humanities Advanced Technology and Information Institute, <http://www.hatii.arts.gla.ac.uk>



University of Glasgow has led the support and coordination activity through a dedicated LOCKSS Technical Support Service (LTSS). Content Complete Ltd⁸, the JISC's negotiation agent for the NESLi2 national e-journals initiative, is undertaking publisher negotiations to obtain permission to allow LOCKSS-based archiving of journal content. Hockx-Yu has described the components of the LOCKSS Pilot Programme (Hockx-Yu, 2006). Throughout this paper we make reference to these components as we detail our progress to date.

The LOCKSS Pilot Programme brings together a diverse set of UK Higher Education libraries. The library budgets are spread across the JISC banding arrangements⁹, and the subscription arrangements differ according to the finances available. E-journals accessible from these institutions vary in quantity, ranging from approximately 7,500 to over 20,000 titles, and these e-journals are made available through many combinations of publishers and aggregators. In addition, the collections include current titles and back-files, and both paid-for and open access journals. The variety of financial, organisational, and structural factors these libraries bring to the UK LOCKSS Pilot Programme supports the exploration of the LOCKSS approach from many perspectives, providing a cross-section that is likely to mirror that which exists globally and across domains.

A programme kick-off meeting was held at the University of Warwick, April 2006. This one-day event introduced UK LOCKSS Pilot Programme participants to the LOCKSS Pilot Programme, its components, its objectives, and our proposed strategy for achieving these goals. The LTSS has held two further events. In October 2006 we delivered a one-day training event covering the operation and maintenance of LOCKSS, and in December 2006 we ran a workshop intended to provide participants with an update on our progress and activities. Participant feedback on these workshops provided evidence of their genuine value to attendees. The workshops also gave us an opportunity to identify themes deserving further attention as the LOCKSS Pilot evolves. The dominant themes which emerge from these sessions are described in detail below, and comprise: content, including identification, diversity, and aggregator access; technology, including hardware requirements, proxy integration, and plugin development; and next steps, including assessment and sustainability. In the discussion that follows we focus on some of the technical challenges and the steps taken to address these. There are other organisational, cultural, and structural considerations that deserve consideration.

Technology: Hardware Environment

The LOCKSS team has identified a number of requirements it believes should influence the design of digital preservation systems (Rosenthal et al, 2005), recommending systems that depend on neither high financial investment nor substantial local technical expertise. The highly automated nature of the LOCKSS platform reduces technical support requirements. For example, the LOCKSS software runs off a dedicated platform CD providing a largely preconfigured secure environment. We were able to reduce the required level of local expertise further by distributing machines conforming to a standard specification, which eased the identification and resolution of problems. For the original 24 participating institutions, the LTSS coordinated a bulk purchase of low-cost computers (£500 per machine in

8 Content Complete Ltd, <http://www.contentcomplete.com/>

9 JISC Collections, http://www.jisc-collections.ac.uk/jisc_banding/collection_banding_he1.aspx



2005). This process was not entirely snag-free; a minor issue we encountered with this approach provides a cautionary note to similar initiatives. The machines we purchased arrived with SATA (Serial Advanced Technology Attachment) hard drives, now distributed as standard with new computers. Our beta machines had not had such a drive. At the time, the most up-to-date version of the OpenBSD operating system upon which the LOCKSS software runs did not natively support SATA. To a certain extent, this is a common issue within the open source community. For very new hardware the open source community, and in particular the smaller OpenBSD community, may not have had an opportunity to develop appropriate drivers. In this instance, the Stanford LOCKSS team developed and released an updated version of the LOCKSS platform containing the required support. The LOCKSS team has consistently taken a measured approach towards platform development and system integrity. The obstacle we encountered and the introduction of an additional, dedicated platform has not caused problems for existing LOCKSS systems. The LOCKSS team supports the two most recent platform releases at any one time, and the manner in which our SATA-support issue was resolved offers a small insight into the benefits of belonging to the LOCKSS Alliance and of collaborative development.

Using the guidelines provided by the LTSS, most participating institutions were able to complete installation without difficulty. LOCKSS requires direct access to the internet on certain ports and where problems were encountered we found these were often caused by either an incorrectly configured network or decisions by particular institutions to lock down port ranges.

Technology: Proxy Integration

Over recent months we have been exploring the proposed methods by which libraries can access the content stored within a LOCKSS box. LOCKSS acts as a transparent HTTP proxy server, with a local cache of the content it preserves. This was designed to be integrated with an institutional proxy, meaning access to content is intended to be transparent to users. In effect a user would perceive no difference between content accessed from an original publisher to that from the LOCKSS box. This can be achieved by configuring a browser or an existing proxy server (such as Squid using the ICP protocol, or EZproxy) to redirect to the LOCKSS box content requests only for those hostnames for which content is known to be stored. This approach is necessary because low-power LOCKSS machines do not have the capacity to handle all institutional HTTP requests. For example, there would be no value in sending to the LOCKSS box requests for <http://www.google.com> as no content for this domain is cached.

The LTSS is actively working with participants to configure their LOCKSS system within their local environment. Feedback from some participants indicates they are hesitant about moving towards a solution where they are reconfiguring their institutional environment for a pilot service. Rather, they are expressing a preference where content in LOCKSS is made available as a distinct resource, more comparable to an institutional repository system, which can be integrated into existing library catalogue software. The US-based LOCKSS team is investigating methods by which this functionality can be added.

***Technology: Plugin Development***

When archiving content, the LOCKSS system harvests an identical copy of journal content from a publisher website which is then stored in the local LOCKSS box. Content is collected in manageable units, often corresponding to subscription units used by publishers. These are termed Archival Units (AUs). Typically, a single archival unit collected by the LOCKSS system matches a complete journal volume.

For each publishing platform, a LOCKSS plugin is used to ensure the LOCKSS daemon correctly collects only the material relevant to a single archival unit. A series of rules corresponding to URLs, known as regular expressions, are used to define the collection of content comprising an archival unit. By harvesting from the publisher's website, the LOCKSS system collects a particular rendition rather than source files. The result is that when users access the archived material the look and feel will match that intended by the original publisher.

As described in the following section, four NESLi2 journal publishers have so far agreed to join the LOCKSS Pilot Programme and we have been undertaking plugin development work in order to release content to Pilot Programme participants. A Java-based Plugin Tool simplifies the construction and subsequent testing. Publishers are first required to create a manifest page for each archival unit (as described above, an AU typically represents a journal volume) to be included within LOCKSS. Each manifest page contains a statement permitting the LOCKSS system to crawl, archive, and serve the content. Since each manifest page is located behind the journal access control list for that particular volume, only LOCKSS boxes at institutions with an active subscription to the particular journal volume are permitted access.

Content: Identification

Content Complete Ltd has been undertaking negotiations with a number of NESLi2 publishers with the objective of securing LOCKSS compliance. We are delighted at the progress made to date: the LTSS is working to release the content of four publishers with which arrangements have been successfully negotiated. The negotiation process can be frustrating at times: conflicting priorities (for example, platform redevelopment and back-issue digitisation) often compete with a commitment to LOCKSS. In addition, the wide variety of emerging digital preservation initiatives has led some publishers to be concerned about committing to too many initiatives.

During the first six months of the pilot programme, commercial publishers remained the primary target of the UK LOCKSS Pilot Programme. The high subscription fees required for the materials supplied by these publishers resulted in a focus by journal preservation initiatives to secure this content, ensuring preservation and post-cancellation access.

In a recent blog post, Rosenthal argues that small, independent open access titles may be at a greater and more immediate risk ([2007](#)). These journals are often published on a shoestring budget and may have no formal management or succession policies in place. The value that these resources bring to an institution may not be formally recognised, and were the managing academic to leave, it is not difficult to imagine a scenario where the title could become unavailable. A straightforward, cost-effective and efficient method needs to be developed by which these publications can be effectively archived. To complement the NESLi2 negotiations undertaken by



Content Complete Ltd, the UK LOCKSS Pilot Programme is targeting open access titles identified as of common interest to the participating libraries.

A survey was distributed to participants in late 2006, with the request that each institution identify ten open access titles they wished to see included in LOCKSS. In March 2006 JISC funded a six-month project based at Glasgow University, OpenLOCKSS (Kidd & Nixon, 2007), with the remit of undertaking negotiations with a prioritised set of open access publishers based on the outcome of the earlier survey. The recent mid-term report (OpenLOCKSS, 2007) highlights the progress made so far. In particular, the project team has focused on developing clear documentation and FAQs which address some common misconceptions, and on streamlining the technical process so that content can be made available with minimum effort by publishers. We have found establishing dialogue especially important, as publishers often had misconceptions as to what was going to happen to their content. Where and how would the content subsequently be accessed? Did its archiving in a LOCKSS environment affect their ownership? Do the publishers formally have the rights from the author to “distribute” their journal content for archiving? These questions continue to be explored by the OpenLOCKSS Project and will be discussed in more detail when they are better understood at the conclusion of OpenLOCKSS.

As these open access publishers commit to LOCKSS (as of June 2007, 13 of the identified publishers have joined the initiative, with another five expressing strong interest), the LTSS is working with the US-based LOCKSS team to improve the documentation describing the plug-in development process (see the following section on Plug-in Development). A key objective for the remainder of the pilot is to disseminate these experiences into the UK community, by distributing improved documentation and delivering focused training sessions.

Content: Diversity

The LOCKSS system is designed to be format-agnostic, meaning it is able to archive all types of digital materials that are transmitted over the web. It is increasingly significant that, as authors become more familiar with multimedia presentation forms and publishers wish to disseminate associated supplementary material such as scientific datasets or software source code and binaries, archiving systems must be designed with sufficient flexibility to handle a diverse range of content types and structures.

Within the LOCKSS Pilot Programme participants have expressed an interest in exploring the use of LOCKSS for a variety of digital content types, including electronic institutional reports, theses, and dissertations. We are in communication with the US-based MetaArchive Project¹⁰, which has been preserving a variety of cultural artifacts for universities in the South-Eastern US. In addition, several titles negotiated through the OpenLOCKSS project make content available through an ePrints repository system. The release of this content will offer an insight into complexities that may arise when using LOCKSS with repository systems.

¹⁰ MetaArchive Project <http://www.metaarchive.org/>

***Content: Journal Aggregation Services***

The LTSS has provided support to familiarise librarians with the daily operation of LOCKSS, and in particular the circumstances under which they were permitted to archive content. Librarians found that often, despite their institution apparently having access to a particular journal, they were unable to collect that content via LOCKSS. This state of affairs reflects how LOCKSS works on business and technical levels. An agreement for LOCKSS compliance is made directly with the original publisher. LOCKSS permission statements are uploaded to the original publisher's website and access is restricted, most commonly, to authenticated IP subnet ranges. With this approach, access and crawling techniques which underpin LOCKSS will only succeed if the institution holds an active subscription to the particular journal. It is an exact copy of the journal from the publisher's website that is collected and preserved in a LOCKSS box.

Many UK institutions subscribe to journals through intermediary aggregators such as Gale or Expanded Academic ASAP. These services effectively republish content with an emphasis on low-cost access. The licence arrangements between the aggregator service and publisher do not extend to cover the use of LOCKSS to archive content republished through intermediary aggregators even where the publisher allows LOCKSS-based archiving. Consequently, although institutions are able to provide users with access to a particular title, they do not have sufficient contractual rights to archive titles accessed through nearly all aggregator services.

For some institutions, this represents a significant proportion of their journal collections. Institutions can only secure perpetual access provisions for their journal content by subscribing directly with the original publishers. Note this restriction not only applies to the LOCKSS approach; designated third-party perpetual access services such as Portico will only make content available to institutions where a subscription has been held with the original publisher. Some pilot participants have indicated that the financial implications of switching their subscriptions from aggregators to publishers requires a strategic change that is not possible in the immediate future, especially where the likely costs of change are high.

The problems described above highlight some of the complexities encountered within the current journal licensing landscape. While providing cost benefits for libraries, the array of licensing conditions, publisher deals, and service providers can make it difficult to understand how content can be suitably used and managed.

Next Steps: Assessment

In the last quarter of 2007, we are making significant increases in the quantities of journal content available to UK Pilot Programme participants. The release of this content follows successful negotiations by Content Complete Ltd. Open access journals identified and negotiated as part of the OpenLOCKSS project are a priority. Many lessons have been learnt from the LOCKSS Pilot Programme's publisher negotiation and plugin development activities, and we will share these experiences with readers in subsequent papers.

One of the benefits of having a large user group of over thirty institutions is our ability to identify gaps in the current system and strategy. Where could the LOCKSS system be strengthened? What has been lacking so far from the UK LOCKSS Pilot



Programme? What advice can we offer institutions with regard to developing their continuing journal archiving strategy? In order to gain an appreciation of library concerns, we are undertaking a comprehensive evaluation of the experiences of the thirty UK LOCKSS Pilot participants and their changing expectations. Following on from this survey, we plan to hold a series of focus groups and workshops to explore in greater detail the issues that emerge during the survey and, where possible, address and resolve them. We hope that this will give us an indication of the future needs of the community for LOCKSS-style archiving.

Finally, we are aware institutions are assessing their medium-term plans for journal archiving. We are currently undertaking cost assessments to establish a method by which the UK LOCKSS Technical Support Service could continue as a sustainable service to provide ongoing support to the use of LOCKSS in UK academic institutions. Continuation of the UK LOCKSS Alliance depends on addressing such issues as the level of LOCKSS Alliance contributions, the levels of staffing required, and the contributions that may be necessary for ongoing publisher negotiations.

Next Steps: Sustainability

The LOCKSS Alliance¹¹, established in 2005, is a membership organisation governed by a Board and advised by a publisher committee, intended to offer institutions a forum to share experiences and concerns related to LOCKSS and journal archiving more generally. Members are offered strategic opportunities to help determine long-term priorities and directions for the evolution of the LOCKSS software and programme. Membership requires an annual fee; the level of this fee reflects institution size and budget. Alliance membership gives member institutions access to premium content, ongoing support, and direct engagement with the LOCKSS development team. The UK LOCKSS Pilot Programme was made possible by a consortium agreement between JISC/CURL and the LOCKSS Alliance.

The current UK LOCKSS Pilot Programme will run until the end of February 2008, but we are putting into place mechanisms by which the UK LOCKSS Programme will achieve sustainability and a life beyond the pilot phase. The central UK coordination has proved valuable by ensuring UK-specific issues are effectively identified and resolved consistently and at national level. By bringing together institutions to share experiences, we are facilitating the development within the information management and library communities of the concepts and issues surrounding journal archiving. This familiarity helps institutions make an informed decision regarding their ongoing journal archiving strategy.

A key priority for ensuring the long-term sustainability of the LOCKSS approach is to ensure that the development expertise, especially that reflecting specific UK priorities, continues to expand. Despite the distributed nature of the LOCKSS technology, the system depends upon coordinated configuration of many aspects and coordinated communication with the US LOCKSS team. The LTSS plays a key role in these activities. The LTSS has contributed to the enhancement of aspects of the LOCKSS system, and has significant familiarity with the content release system. The continued viability of LOCKSS depends on the efforts and success of initiatives led by the LOCKSS foundation and development team in the US. Other national initiatives, following the model used by the UK LOCKSS Pilot Programme, will result in a wider

11 LOCKSS Alliance http://www.lockss.org/lockss/LOCKSS_Alliance



distribution of development knowledge. In addition, local responsibility will ensure that the best selection of at-risk journal content is made and archiving agreements are further negotiated.

Conclusions

In response to the concerns of librarians, publishers are increasingly participating in efforts to develop effective journal archiving strategies. The LOCKSS system provides a critical component in the journal distribution infrastructure, allowing libraries to take custody of assets for which they have paid, while conforming to the licensing arrangements they have agreed with publishers. The LOCKSS approach takes steps to ensure libraries are responsible not only for short-term access, but involved at many stages in the emerging model of journal archiving. The UK LOCKSS Pilot Programme demonstrated a way in which an effective LOCKSS community can be established and run. It provides a model for other national, regional, or trans-institutional approaches. It is certainly the case that the technologies and licensing agreements will continue to develop and evolve, ensuring that both libraries and publishers acquire, secure, and maintain the rights, access conditions, and financial benefits that they both expect and deserve. The experiences, so far, of the UK LOCKSS Pilot are that the future access to academic journals depends not only on technical solutions but also on ensuring that cultural, financial, and organisational components and activities are adequately aligned.

Acknowledgements

We would like to thank Helen Hockx-Yu, The British Library, formerly of the Joint Information Systems Committee (JISC), for her comments and suggestions in respect of this article. The Stanford LOCKSS team and in particular Vicky Reich and David Rosenthal have provided invaluable assistance and guidance on many occasions. The UK LOCKSS Pilot Programme would not have been possible without the participation and support of Content Complete Ltd, the 30 institutions, and the invaluable feedback offered by each.

References

- Boulderstone, R. (2007). *The British Library e-Journal Digital Archive*. Powerpoint presentation at E-Journal Archiving and Preservation Workshop, March 27, 2007. Retrieved December 5, 2007, from <http://www.dpconline.org/docs/events/0703ejournalboulderstone.pdf>
- Fenton, E., (2006, April) Preserving electronic scholarly journals: Portico. *Ariadne Issue 47*. Retrieved December 5, 2007, from <http://www.ariadne.ac.uk/issue47/fenton/>
- Hockx-Yu, H. (2006). Establishing a UK LOCKSS Pilot Programme. *Serials: The Journal for the Serials Community*, Volume 19 Number 1. Retrieved December 5, 2007 from <http://uksg.metapress.com/openurl.asp?genre=article&issn=0953-0460&volume=19&issue=1&spage=47>

- 
- Jones, M. (2007). *Review and Analysis of the CLIR Report "E-Journal Archiving Metes and Bounds: A Survey of the Landscape"*. Retrieved December 5, 2007, from http://www.jisc.ac.uk/whatwedo/programmes/programme_preservation/ejournalarchiving.aspx
- Kenney, A., Entlich, R., Hirtle, P., McGovern, N., & Buckley, E. (2006). *E-Journal Archiving Metes and Bounds: A Survey of the Landscape*. Council on Library and Information Resources Report. 120 pp. Retrieved December 5, 2007, from <http://www.clir.org/pubs/abstract/pub138abst.html>
- Kidd, T., & Nixon, W. (2007). Proposal to JISC from the University of Glasgow for negotiations with UK Open Access Publishers. Retrieved December 5, 2007, from <http://www.lib.gla.ac.uk/Research/openlockss/jisclockssproposalpublic.pdf>
- Maniatis, P., Roussopoulos, M., Giuli, T.J., Rosenthal, D., Baker, M., & Muliadi, (2003). Preserving peer replicas by rate-limited sampled voting. In *19th ACM Symposium on Operating Systems Principles (SOSP)*. Retrieved December 5, 2007, from <http://www.eecs.harvard.edu/~mema/publications/SOSP2003.pdf>
- Oltmans, E., & van Wijngaarden, H., (2006). The KB e-Depot digital archiving policy. *Library Hi Tech* 24(4):pp.604-613. Retrieved December 5, 2007, from <http://eprints.rclis.org/archive/00009159/>
- OpenLOCKSS Team. (2007). *OpenLOCKSS mid-term project report*. Retrieved December 5, 2007, from <http://www.lib.gla.ac.uk/Research/openlockss/openlockssreportmidterms.pdf>
- Reich, V. (2006, April). CLOCKSS: Controlled Lots of Copies Keep Stuff Safe – Large dark archive initiative. *D-Lib Magazine, Volume 12 Number 4*. Retrieved December 5, 2007, from <http://www.dlib.org/dlib/april06/04inbrief.html#REICH>
- Reich, V., & Rosenthal, D. (2001, June). LOCKSS: A permanent web publishing and access system. *D-Lib Magazine, Volume 7 Number 6*. Retrieved December 5, 2007, from <http://www.dlib.org/dlib/june01/reich/06reich.html>
- Rosenthal, D. (2003). A digital preservation network appliance based on OpenBSD. In *Proceedings of BSDcon 2003*, San Mateo, CA, USA, September 2003,. Retrieved December 5, 2007 from http://www.usenix.org/publications/library/proceedings/bsdcon03/tech/full_papers/rosenthal/rosenthal.pdf
- Rosenthal, D. (2007). Why Preserve E-Journals? To Preserve The Record. Posting on personal weblog, <http://blog.dshr.org>, June 10, 2007. Retrieved December 5, 2007, from <http://blog.dshr.org/2007/06/why-preserve-e-journals-to-preserve.html>



Rosenthal, D., Robertson, T., Lipkis, T., Reich, V., & Morabito, S. (2005, November). Requirements for digital preservation systems: A bottom-up approach, *D-Lib Magazine, Volume 11 Number 11*. Retrieved December 5, 2007, from <http://www.dlib.org/dlib/november05/rosenthal/11rosenthal.html>

Waters, Donald J., et al. (2005). *Urgent Action Needed to Preserve Scholarly Electronic Journals*. Retrieved December 5, 2007, from <http://www.diglib.org/pubs/waters051015.pdf>