The dogs bark and the circus moves on

John Robinson
Director of Library & Information Services
SOAS, University of London
January 2015

The dogs bark and the circus moves on

In 2008 I wrote a chapter for *Digital Library Economics*[1] called "Spinning the disks: lessons from the circus"[2] which started from the proposition that it is feasible and useful to attempt to describe the infrastructure required by the digital library in terms which can be understood by a non-technical reader. I used the metaphor of the master practitioner of plate-spinning to describe the IT practice of keeping data on "spinning disks" and went into some detail to set out how a costing model for a Digital Library (by which I meant a traditional print-based library moving towards a hybrid of print and electronic content) could be developed. I concluded by suggesting that, from the point of view of a library that has traditionally seen itself as the collector, curator and protector of printed materials for the long-term, information held in purely digital form is by its nature ephemeral. The threat of a "digital dark age" is a profound epistemological challenge.

Writing now in 2015 I would suggest that this challenge is, if anything, greater, despite the many efforts by individuals and institutions to find solutions. I will touch briefly upon the question of sustainability in relation to electronic resource subscriptions before dealing with the new challenges for the Digital Library presented by the Open Access publication and Research Data Management mandates.

I should say that the distinction between the Library (as we think of it in the Higher Education Context) and the Digital Library is used to foreground the things which are new or which change with the emergence of the digital realm, but the distinction is increasingly meaningless as all libraries (with rare exceptions) are now hybrids of digital and print media – even if the digital is restricted to the methods by which the print materials are catalogued, described and discovered by the library users.

Any UK HEI that wishes to participate in the 2020 Research Excellence Framework (REF) – or its equivalent – has been put on notice that the majority of research outputs for inclusion in this exercise must have been published on an Open Access basis from 2016 at the latest. Many research funders are now making it a requirement of their grants that not only the research outputs, but also the research data, must be preserved and made available – preferably on an Open Access basis – to other researchers.

These new obligations fall in many cases on libraries as resource managers if not actually resource providers. This article will examine the service models that support these new obligations and how their costs might be estimated.

The Digital Library as subscriber

At the time of writing in 2008, I focused on the issues (and new costs) that had arisen from the move to electronic publication of journals. In 2015, despite concerted campaigns by organisations such as RLUK[3] and SCONUL[4], ably assisted by Jisc Collections[5], the costs of journal subscriptions continue to increase well above inflation whilst university budgets are squeezed, leading to difficult choices for subscription librarians. Electronic books in the academic context are still a relatively new phenomenon with their own sets of challenges

(and cost models) but they are already putting pressure on the traditional book fund and challenging the notion of "ownership" of books.

The problem of long-term access continues to worry us.   Accurate and reliable citation of journal sources remains a pre-requisite for scholarly writing, and it is a fundamental assumption that the reader may, if she chooses, go to the original source to check the citation, perhaps to undertake further research.

What happens when the location of the journal in question is "on-line at the publisher"?   What if our subscription has lapsed?   What if the publisher has gone out of business and closed down its on-line systems?   What, in those cases, has happened to all the material for which we have paid our subscription charges for all those years?

The notion of "escrow" as used in the world of computing (source code placed in the hands of a secure third party in case a supplier goes out of business), has not yet become pervasive in the world of electronic journals.   There are various solutions such as Jstor[6], LOCKKS[7], CLOCKSS[8], Portico[9] and the BL Document Supply Service[10], but as yet no comprehensive, universal solution to this problem.

One thing though is certain: paying a third party to hold secure copies of the journals to which we subscribe is an extra cost.   The other thing which is clear from the examples cited is that the most effective methods for dealing with this challenge are collaborative, collective or centralised solutions.   These methods provide us with examples of the models we may wish to adopt when dealing with our new digital curation and preservation obligations.

The Digital Library as repository

Libraries have traditionally dealt with holdings: the books and journals that we collect and hold for the benefits of our owners, our members, our users or (for Research Libraries), posterity.   Holdings are straightforward to manage: we buy them (or they are given to us), we catalogue them, we mark them up, we put them on our shelves or into our closed reserves for reference or borrowing, we keep track of them, we relegate the lesser-used materials to store and weed those we no longer need.   Subscriptions, in this model, are the recurring payments for serial publications.   The journals that arrive in printed form are eventually bound into volumes, which we can hold on our shelves as easily as monographs.   We also already hold a certain amount of the intellectual property of our institutions such as the theses presented for assessment by our Research Postgraduates.

A Digital Library (especially in the Research Library context, but realistically any library in a UK HEI which subscribes to the Research Excellence Framework) is in many cases becoming the primary repository for its host institution's intellectual property (with good reason).   These are new types of holding for the library but the sense of their importance to the institution is familiar: any library which has a large Archives, Rare Books and Manuscripts section already knows how to preserve valuable, irreplaceable materials.   If the library also holds its institution's records, it knows about the concept of "vital records" (those records upon which the life of the institution depends).   Digital curation and preservation

is not a new concept; what is new is the way that the Open Access and Research Data Management mandates have put these obligations into sharp focus.

The Digital Library as publisher

Open Access is not a new concept.   There have been a number of initiatives over the past decade, such as "Securing a Hybrid Environment for Research Preservation and Access" (SHERPA)[11], "the London E-prints Access Project" (SHERPA-LEAP)[12].   Many other projects pop up under the SHERPA link including

- RoMEO - Publisher's copyright & archiving policies
- JULIET - Research funders archiving mandates and guidelines
- OpenDOAR - worldwide Directory of Open Access Repositories
- SHERPA Search - simple full-text search of UK repositories

revealing, through their interesting acronyms, their origins in the Jisc project funding which lay behind them (at one time the provision of an acronym was mandatory for Jisc funding).

What has changed is the movement in the last year or two to make Open Access a mandatory requirement for funded research.   The Wellcome Trust led the field in this regard, with its requirement that all its funded research must be put on open access.[13]   There followed the HEFCE[14] and RCUK[15] policies.   On the HEFCE page, Libraries are a key part of its diagram, showing the "Green" and "Gold" routes to Open Access publication.

The reason why "Green" is of interest to libraries is reasonably clear, given the involvement of libraries (alongside colleagues in IT and research departments) in the various initiatives which have resulted in "research repositories" becoming in many cases (but not all cf "discipline-based repositories"), library responsibilities. Green means that a research paper is either published immediately in an open access repository, or placed with a commercial publisher with an agreement that it can be put on open access after a period of time (the "embargo" period). Regardless of the embargo, good practice is for the researcher to put the bibliographic details of the publication immediately into the institutional repository so that it can be indexed and exposed via the institution's web pages and the library's discovery systems.

For libraries, the "Gold" route (payment up-front of an "Article Processing Charge" so that the article is made freely-available by the publisher) is more complex.   On the one hand, the published article can be deposited immediately in the institutional repository but its presence in commercial journals to which the library likely subscribes presents an interesting problem.   APCs have proven a gold mine for many publishers and a new cost upon the sector but do libraries see any evidence that subscription costs are reduced as a consequence?

Building the Digital Library service model

As long as libraries have existed, librarians have focused on the challenges of collecting materials and keeping them safe.   This is a task shared by all libraries, regardless of size.   For the most part it has been about the need for secure buildings and methods for logging and tracking the materials.   Research or Special Collections libraries or Archives face a more elaborate version of the

challenge with the need for extra security (in the case of valuable collections), and temperature- and humidity-controlled environments (for fragile materials). These are well-known infrastructure challenges supported by decades or even centuries of practical experience. Successful libraries work closely with their Estates and Facilities departments to ensure the physical challenges are met.

It is not unusual (indeed it is expected) to find materials in libraries, archives or special collections that have been kept and preserved for hundreds of years. Books, journals and manuscripts using good quality paper, held in the right conditions, can be extremely stable; photographs, films and (physical) audio recordings are more problematic but can be preserved for long periods of time. There is a trend, however, in more recent formats such as audio tape, video tape, laser disks, floppy disks, CDs and DVDs where conservation becomes increasingly difficult as the shelf-life of the materials is limited, and the devices required to access the contents increasingly difficult to maintain. This trend reaches its peak with the phenomenon of material that is "born digital". Here we enter the world of the Digital Library by which we mean, the library of stuff that starts, and remains, digital.

Service models for the Digital Library fall into two main components, which are at least familiar to any librarian: infrastructure and support. Major new costs arise from each component.

Digital Library infrastructure

The first question when working up an infrastructure model is to determine whether we are seeking to emulate in digital form the traditional library methods for collecting, holding and preserving material, then regulating its use. If that is the case, the library will want to have direct control over its holdings.

Libraries that seek to be stand-alone digital libraries must take on board new facilities and the skilled staff required to build and to operate them. The extent to which this is feasible will depend upon the precise circumstances – and resources – of the library. A very large library such as the British Library will likely already have investments in IT infrastructure and specialist staff to support this approach; a wealthy library may be able to afford to develop its own IT department.

Much more commonly (in the education sector at least), libraries are required to engage with and depend upon the infrastructure provided by their Computing Service or Information Systems departments

In each case, the generic challenges are the same: capital-intensive infrastructure which must be replaced every five or ten years; teams of people with the high-level systems-administration and programming skills required to operate the infrastructure; new approaches to revenue budgeting which takes into account the new operating costs and the requirement by the institution to defend these costs against its core objectives in learning, teaching and research.

Data Centre and Computing infrastructure (which is what a Digital Library requires) is built up in layers, analogous to constructing a building. If one were to start from scratch to build the infrastructure for a Digital Library, the components required are

- A secure building in which to locate the machine-room

- A specialised environment with qualities similar to a secure archive store such as temperature- and humidity control, but having to cope with hardware which generates a lot of heat

- Secure electrical supply and good network connections to the Library and the outside world

- Large-scale storage arrays with capacity measured in Terabytes (thousands of Gigabytes)

- File servers to manage the storage

- Database servers to manage access to the storage

- Web servers and other specialised systems to index the data and present it to the user

- Systems for making secure backup copies of all vital data

This latter point is worth considering.   No digital data can be considered safe if there are not at least three complete copies held – ideally – in three separate locations.   How do we achieve this?   If we are a truly standalone library with unlimited resources, we will need to build machine rooms in two other locations.   Maybe if we have multiple libraries spread around the landscape, this is feasible.   Otherwise, we are looking for premises to buy or rent, with good network connections back to our main site.

For someone who has run large IT services for large universities, none of this is particularly unusual.   For a library trying to build from scratch, the starting point for calculating costs is

- Purchase or lease a building: unknown, could be millions (to purchase) or hundreds of thousands per annum (to lease)

- Build and equip a machine room: £1M (small room) and up

These first two costs are in the realm of major capital investments, which might seem affordable if one were already building, say, a new library for £20-50M or renovating an existing library for £10-20M.   Adding £1-2M to the cost to include a secure library machine room, perhaps alongside the state-of-the-art Archives store would make sense.   Otherwise, it might be a difficult ask.   The lifetime of capital investment in a building is usually calculated in decades.

Other costs are easier to calculate.   Equipping the facility is a matter of spending 100s of 1,000s of pounds, depending on scale.   For example

- Purchase 150 Terabytes of data storage: £180,000

- Backup system for the data: £125,000

- Network hardware: £50,000

- Server cluster: £50,000

- Incidentals: £75,000

In summary, for about a half million pounds of capital investment (over and above the cost of setting up the facility), one can get into the business of running a Digital Library on-site, where it can be seen, with the data held close, just like any other library facility.

Two things are missing from this calculation:

- The operating costs
- The capital depreciation

For the purpose of the exercise, let's assume we need an IT team to run this (we'll look at the support team later).   Using a full-economic-cost model (which covers salary, superannuation, NI, estates and indirect costs), based on the UK HE National Pay Spine where G9 is the top scale and G3 the lowest, we can build a team.

- IT Manager @ G9 FeC: £82,000
- 2 x Snr Programmers @ G8 FeC: £140,000
- 2 x Support Officers @ G7 FeC: £120,000

That's £342,000 p.a. to add to our staffing costs.   On top of that we need to account for the costs of electricity and maintenance which we'll call £100,000 p.a. and depreciation on the half million capital @ 20% which is another £100,000 p.a.

So, we have a charge on our revenue budget of £542,000 p.a. to keep this facility operating.   There are still two things missing: the two other places to which the data has to be backed-up if it is to be considered truly secure; the team which will support the users of this service.   Back to that later.

If you are not running a major library with an annual budget already running to eight or more figures (or the big end of seven), you are likely wondering: what is the point?   If you are running a library in a medium to large university then these kinds of calculations are a good example of the way in which your IT director (assuming you go to her for assistance) will be thinking.   If you are a small library, the chances are you may be looking elsewhere for your Digital Library facilities.

Here be dragons: fat pipes, clouds, snake-oil vendors, shared service providers

In the UK HE sector and most of the FE sector, we are blessed with "fat pipes" into something called "Janet".   Janet recently merged with Jisc (which used to be its funder), and they now offer a number of services that can help smaller libraries.   For facilities, there are these "fat pipes" into shared data centres that have "big racks" full of "tin" (to paraphrase some of the dialog which was overheard at a recent meeting on the topic).   Other service providers are also available.

The notion of "cloud computing" has become ubiquitous in recent times.   It is worth thinking through the origins of the term.   It originated with computer network engineers faced with the difficulties of drawing accurate diagrams of their networks.   There is no simple way to describe a computer network: there can be a physical view, a logical view, an organisational view, a functional view.   The "cloud" in network diagrams was devised as a means to show a section of the network to which one has a connection, or through which data might traverse, without having to show any details of what lies within.   The use of the

term "cloud" is obfuscation: a function is provided to a user without the need to reveal any detail of how it operates.

In the realm of digital services, cloud has become a marketing term: subscribe to, or purchase, or simply use (in the case of "free") a service whilst allowing someone else to worry about the details.   In fact, the term is used to describe any number of different services ranging from externally-managed infrastructure, through "software as a service" and externally-managed content to discrete applications, all delivered across, and reliant upon, continuous and ubiquitous network connections (those "fat pipes").

For a librarian, the key point when offered an opportunity to "put it in the cloud" is to understand that what this really means is "put your data on someone else's computer".   Another way of stating this is, "put your data beyond your direct control".   The matter of trust and confidence is an essential consideration: "do I trust this external supplier to whom I am outsourcing my obligation to protect and preserve my data?"   To some extent, the answer will depend upon what kind of data I am proposing to send off-site.   Is it my most rare and valuable data (e.g. primary research data), is it my confidential data (e.g. the personal data of my users), or is it data that, were it to be lost, I know I can easily replace from a different source?

There are a number of health warnings about outsourcing the Digital Library infrastructure.   One thing which has changed dramatically since 2008 is the revelation and widespread understanding of something which IT professionals have known for some time: that any and all material held on line, in data centres, accessed via computer networks, can be accessed not only by those professionals who manage the facilities but also by government "security" agencies and other more obscure agents (who may have commercial or criminal motives).   What might have seemed a "safe" option in 2008 (digitising material and putting it on-line for reasons of preservation and accessibility) is now safe only if one is not overly concerned with privacy or confidentiality.

The second is more of a practical consideration: if I put my data somewhere else, what happens if I want to get that data back?   Say I put my 150 Terabytes data on an outsourced provider but the provider goes bust, or puts its prices up, or I simply don't like its service any more: how do I get it back?

This touches (again) on the "three copies" rule.   If the provider goes bust, and assuming I can purchase a local 150 Terabyte storage array to restore the data locally, where is my safety copy?

Even if I have opted for local primary storage, I still need those remote copies.   Whatever my arrangements, if I can't afford to run my own facilities in three separate locations, chances are I will be dealing with a shared service or an outsourced service – or even that mysterious "cloud" – for my backups.

Digital Library support services

Once the facility question has been addressed, and assuming we now have a reliable Digital Library infrastructure (which includes capacity to secure the original research data), we need to support the users.   We are back in much more conventional staffing territory here.   Who are our users, what sort of

support do they need?   What are the specialist skills that are required?   Do we add these tasks to existing teams, or start new teams?

If funding is available, starting a new team is tempting.   A number of UK Research Libraries have piloted various approaches – in some cases with Jisc funding – and it is not unusual now to find new jobs in Digital Curation or Research Data Management springing up, in teams such as "RDM Support" or "Special Collections and Digital Archives".   Another question is how many of the existing library services – the Library Management System (LMS), the Discovery System (OPAC), the Archives system – to combine into the team.   If there is already a strong Library Systems team in place, this might be a good place to start.

A model Library Digital Services might look like this:

- Digital Services Manager @ G9 FeC: £82,000
- 2 x Library Systems Analysts @ G8 FeC: £140,000
- Digital Archivist @ G7 FeC: £60,000
- Training and Support Officer @ G7 FeC: £60,000
- Digitisation Assistant @ G5 FeC: £39,500

Within a Library Senior Management structure, this team might report to a Head of Digital and Special Collections, or a Deputy Librarian, or Head of Collection Management or Associate Director (depending upon how the Library is organised).   In a small Library, it might report directly to the Librarian (and might be a considerably smaller team).

The primary tasks allocated to these teams might include

- Provide expert advice in the adoption and use of library systems (depending upon the scope of the team)
- Liaise with suppliers of library, open access and research data management systems to ensure that functional requirements are understood, service levels are met, upgrades planned and implemented
- Support the users of the systems in cases of technical problems or with the loading of data
- Deal with copyright issues in relation to material being put on-line
- Provide training, particularly where systems are new or new users (e.g. early-career researchers) need to be helped to develop good working practices
- Other activities within the scope of the team's remit

There is a clear challenge to any library of any size if its ability to grow its staffing base is constrained, to identify staff members who can be retrained to take on these new tasks and to review its activities to identify tasks that can be modified or dropped to make way for the new.   In some cases, technology may assist.   If staff are spending a great deal of time processing book invoices by hand, an upgraded library system might turn this into an automated process. Cataloguing staff may benefit from access to on-line bibliographic records and the purchase of shelf-ready stock.   Text books may be available as e-books,

reducing the annual pressure to order new stock and weed the stock which is no longer required.

For any library seeking to modernise, these new responsibilities have a direct impact on the career aspirations of professional library staff. New roles of data librarian or open access publications officer require new skills. Staff coming in may have these skills; existing staff can benefit from targeted CPD (see below). Subject or Academic Support librarians can be persuaded of the benefits of adding Open Access publications and Research Data Management support skills to their CVs.

The other challenge – which in some libraries may be quite a big one – is for library and IT staff to work much more closely in collaborative or partnership arrangements as libraries move into these new technologies. In Arts, Humanities and Social Science institutions, it may well be that libraries (alongside learning support) are the most active users of new and innovative technologies; where there are large Computer Science or Engineering schools, it could be that the programmers who are writing the most interesting software (frequently using open-source methods) are right on campus.

Cost summary

If we take as a model, a standalone library which wants to build its Digital Library capability from scratch, in order to become the provide of Open Access publication and Research Data Management, we can provide a cost summary which looks like this.

We will assume for the purposes of this exercise that the library owns premises with sufficient space to build the facilities it needs and that it is a medium-sized library with an annual operating budget of around £10M.

1. Capital investment (building and equipment)   £2.5M
2. Recurrent facility costs (IT Team, overheads, depreciation)   £0.54M
3. Recurrent support costs (library OA & RDM team)   £0.4M

These figures can be used to benchmark a service approach. A capital investment is required, plus an increase in operating budget of around 10%.

Where a library cannot raise capital and buys in to a managed service for its facilities (either internal, from the IT department, or from an external provider), the facility cost is a useful starting point to set a budget for the offer.

If the library can't make the case for new staff to form its OA & RDM team, the challenge is to repurpose existing posts.

Library as publisher

The library as custodian of its institution's repositories for intellectual property has taken one step towards a re-invention of the library as publisher. Institutions may already hold a monograph imprint and publish journals. Many of the great names in publishing originated as university presses. Practice in recent decades has swung towards publication via commercial presses, and there are those in the sector that see this as an unreasonable arrangement in which profits are generated through the exploitation of publicly-funded research which has been peer-reviewed by academics for little or no remuneration.

When publication required large investments in printing and distribution infrastructures, the use of commercial publishers had a certain logic to it, especially when institutions decided – using management jargon popular in the 1990s and 2000s – that running large printing presses was not "core business". New technology has changed the situation dramatically: digital curation and preservation *is* core business.

The spectre of e-journals, in which publishers raise prices every year well above the rate of inflation but deliver nothing other than a subscription to an on-line system which ceases if the subscription is cancelled (pushing libraries towards expensive add-on services to protect long-term access to back-files) has brought this situation into sharp focus.   E-books are a more recent phenomenon.   Given that institutions have invested in repository technologies, the case against self-publishing has become much weaker, and institutions are starting to take an interest in taking back their publication rights: firstly for monographs (a number of imprints have recently sprung back into life); secondly for journals.

What is the difference between "vanity publishing" and scholarly publication that will be taken seriously?   Put simply: peer review.   Libraries can take a leading role in managing their institutions' titles and imprints.   What the institutions need to do is form editorial boards to oversee the publications, validated by strong methods of peer review.

This goes beyond the remit of libraries acting alone.   Examples already exist to demonstrate the feasibility of such an approach[16].   For smaller institutions the best way forward may be to work collaboratively on shared presses.   Ubiqity Press[17] is a well-known example of this approach.   Other publishers are also available.

At the rate of change in this environment it might seem foolhardy to make predictions but perhaps it is not beyond the realms of fantasy to imagine that, 6 or 10 years from now, the spectre of commercial publishers running 35% profit margins on the back of scholarly communications will be nothing but an historical blip in the history of academic publishing.

Research data management

A great deal of work has been done to develop models of good practice in this field.   RDMRose at the University of Sheffield is one of the best-known, "a JISC funded project to produce taught and continuing professional development (CPD) learning materials in Research Data Management (RDM) tailored for Information professionals".[18]   The Digital Curation Centre in Edinburgh has also prepared RDM Training for librarians.[19]   The White Rose consortium of libraries has published a report into "research carried out at the end of 2012 to survey UK universities to understand in detail the ways in which libraries are currently involved in research data management (RDM) and the extent to which the development of RDM services is a strategic priority for them".[20]

The focus here is on what libraries can do.   We have mentioned the costs of infrastructure with an example of 150 Terabytes of file store and the problem of finding ways to ensure that there are at least three copies of all vital data (which includes primary research data) in three different places.   For libraries based in large research institutions that deal with the physical sciences (generally known as STEM), there will be colleagues in IT or large departments who have been

grappling with this for some time, and may have been dealing in Petabytes rather than Terabytes.   Many research councils have funded discipline-specific repositories such as the ESRC-funded UK Data Archive at the University of Essex[21].

For Arts and Humanities the situation has been more complex since the AHRC closed down its discipline-specific Arts and Humanities Data Service (AHDS)[22] in the late-2000s, declaring that long-term preservation was an institutional responsibility.   Whilst data volumes have been lower in these fields, the situation is changing rapidly as researchers are starting to use technologies such as high-definition video for their primary data.   This burden can be a shock to institutions that have not had to invest heavily in large-scale data storage until recently.

Libraries can bring two things to this field: deep knowledge of metadata, cataloguing and indexing taxonomies (often lacking in the methods by which large volumes of data have been collected); experience in supporting postgraduate and career researchers in their methods, which must now extend to having robust data management plans.

The issues are now well known and there is much evidence of good practice to build upon.   For libraries, the biggest challenge can be simply finding the staff – or the funding for staff – to build the new support teams which are required.

Conclusion: keeping the spinning disks aloft

I started in 2008 with the image of the circus performer, dashing about the stage, keeping plates spinning at the end of bamboo poles.   When done well, this is a masterly performance, rightfully deserving of respect.   When considering the digital library, the metaphor has resonance in the first instance because of the notion that the data which the library contains is held on spinning computer disks in order to be delivered on-line, upon demand.   This performance depends upon the efforts of highly trained IT professionals working with expensive facilities that were, in 2008, beyond the experience and expertise of many libraries.   The spectre I raised was the epistemological challenge of access to information being mediated by "invisible hands" outside the domain of the librarian, which had been for centuries, the primary gatekeeper to the repositories in which the information resided.

In 2015 we are in an environment where many minds have been exercised by this problem and many solutions devised.   By and large, the theme in the solutions has been cooperation and collaboration around the shared ethos of academic freedom enabled by open and accessible scholarly communications. This is an approach in which libraries excel, and an ethos that has informed them at least since the time of the Enlightenment if not before.

The challenge for libraries which are now reinventing portions of themselves as Digital Libraries (never forgetting that print is still the most sustainable medium in the long-term, not about to go away) is to regain their primary function as the custodians of the intellectual property of their organisations, as evidenced by becoming their Publishers (through the Open Access mandate) and their Digital Curators (by taking on the challenge to ensure that primary research data is well-managed and preserved for the long-term).

The spinning disks are bigger, the challenges greater.   For librarians, it has to be said, we are living in very interesting times.

John Robinson, Director of Library & Information Services
SOAS, University of London
January 2015

---

[1] Robinson, John (2009) 'Spinning the disks - lessons from the circus.'   In Baker, David and Evans, Wendy, (eds.), Digital Library Economics.   Oxford: Chandos Publishing, pp. 227-246. (Chandos Information Professional Series)

[2] http://eprints.soas.ac.uk/12429/

[3] http://www.rluk.ac.uk

[4] http://www.sconul.ac.uk

[5] http://www.jisc-collections.ac.uk

[6] http://www.jstor.org/

[7] http://www.lockss.org/

[8] http://www.clockss.org/clockss/Home

[9] http://www.portico.org/digital-preservation/

[10] http://www.bl.uk/bldss

[11] http://www.sherpa.ac.uk/

[12] http://www.sherpa-leap.ac.uk/

[13] http://www.wellcome.ac.uk/About-us/Policy/Policy-and-position-statements/WTD002766.htm

[14] http://www.hefce.ac.uk/whatwedo/rsrch/rinfrastruct/oa/oa/

[15] http://www.rcuk.ac.uk/research/openaccess/

[16] see for example http://www.openoasis.org/index.php?option=com_content&view=article&id=348&Itemid=381 and also http://www.manchesteruniversitypress.co.uk/cgi-bin/scribe?showgroup=ig06

[17] http://www.ubiquitypress.com/

[18] http://www.sheffield.ac.uk/is/research/projects/rdmrose

[19] http://www.dcc.ac.uk/resources/developing-rdm-services/rdm-training-librarians

[20] http://eprints.whiterose.ac.uk/76107/

[21] http://www.essex.ac.uk/depts/ukda.aspx

[22] http://www.ahds.ac.uk/