

Main Articles

Implementing Ex Libris's Primo at the University of East Anglia

[Nicholas Lewis](#) outlines the University of East Anglia's experience of implementing Ex Libris's Primo, a new search and retrieval interface for presenting the library catalogue and institutional databases and e-resources.

Main Contents	Section Menu	Email Ariadne	Search Ariadne
-------------------------------	------------------------------	-------------------------------	--------------------------------

Introduction

At the University of East Anglia (UEA), we have been taking part in the Primo Charter Programme in which various libraries in the Europe and the US have been able to work with Ex Libris on version 1 of their Primo product.

We have learned a great deal from the process and there is interest throughout the library sector in the potential benefits of separating or decoupling the search and retrieval interface layer from the database layer when presenting library resources. The problem Primo looks to solve is the overwhelming evidence that users are preferring other search, retrieval and delivery services over the library catalogue and subscribed institutional resources. The reasons for preferring search engines like Google include ease and speed of use, relevance-ranked results, few authentication challenges, and what can be described as the overall aesthetic experience. There are several papers on these issues, but one of the most influential has been OCLC's 2005 report 'Perceptions of Libraries and Information Resources' [1] which has helped to corroborate Ex Libris's strategy. In theory there is little argument that our local services should be developing more in line with the wider network of services such as Google, Amazon and Facebook, but there is still considerable debate as to how it should be done.

Many people argue that a more up-to-date kind of search interface is long overdue in Library catalogues and that Library Management System (LMS) suppliers ought to have provided such functionality as standard as their OPACs (Online Public Access Catalogues) have developed. The counter-argument contends that merely updating the library catalogue is investing a great deal of effort in old technology. To do this well, it is worth investing in a new and flexible search layer that can include resources and databases beyond the library catalogue and which is free from the constraints of the backroom aspects of the Library Management System.

There is another aspect to this debate: there is the argument that since the vast majority of users are now using sites like Google and Facebook, there is little point trying to develop our own local interfaces to resources. Instead we ought to concentrate our efforts on making our resources visible within the wider network of search engines, union catalogues and open repositories. The idea is to concentrate on making the library's local records more visible in the wider network and then to funnel users back to the local resources.

Ex Libris takes such concerns into account but firmly places itself in the ideological camp that claims local is best. Whilst Google and Facebook may have the upper hand in terms of attracting users, they are nonetheless different services. A suitable strategy would be to examine what makes them successful and apply some of their features to our more local interfaces. In so doing we are more likely to encourage people to use our interfaces and as a consequence obtain the highest-quality resources they need to support their teaching and learning. That is where libraries have the upper hand: if content is king, then libraries win hands down due to the quality of the content they have to offer in comparison with Google. The problem lies in the fact that not all our users know this; even if they do, the presentation of those resources via different databases and interfaces often seems less attractive and engaging than what can be found on the Web. This is where Primo steps in. Primo is about finding, not just searching, and delivering resources no matter where they originate. In this respect, Primo is more than just a glorified replacement for the library catalogue.

Front of House

The key benefit of Primo is that it represents a one-stop shop for access to local institutions' library catalogue, institutional repository and remote databases. The relevancy ranking algorithm is sophisticated, leading to results by relevance not just year of publication; it works more like the search engines our users are familiar with on the Web. The visibility of book covers (from Amazon [2]) and icons, rather than text, to distinguish between different types of records, contribute to an enhanced and enriched aesthetic experience.

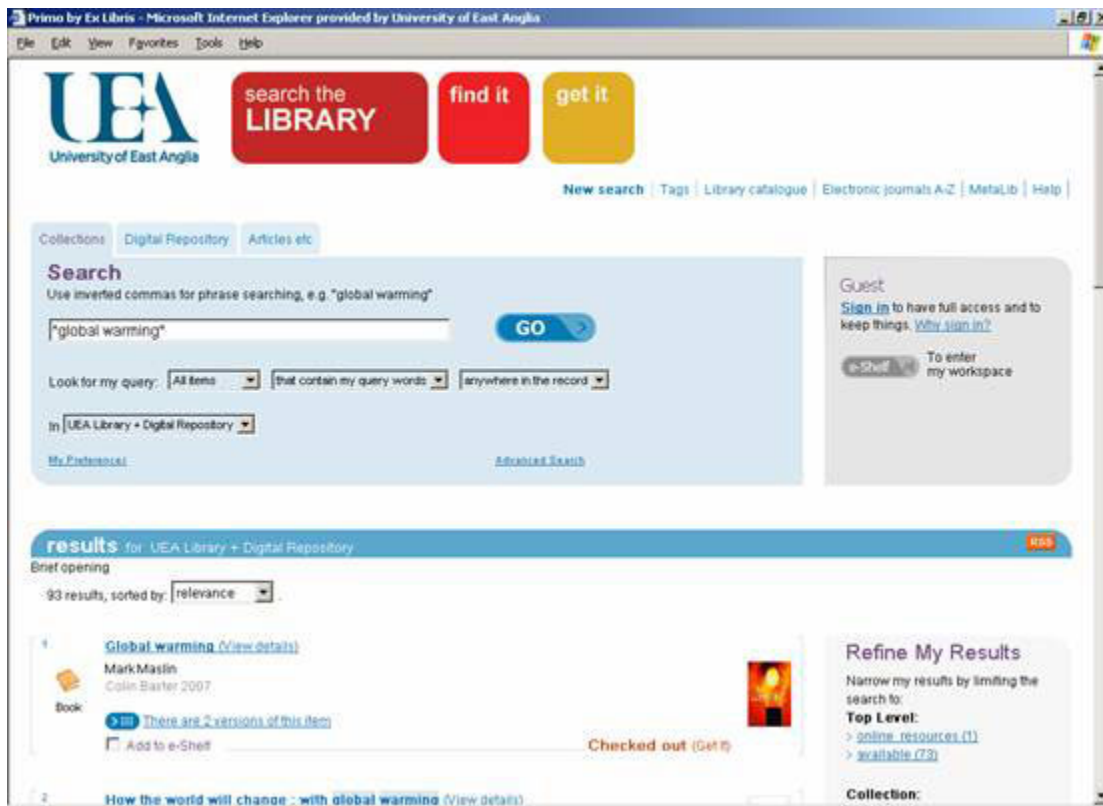


Figure 1: The PRIMO user interface

As well as the main search box, Primo gives users pre-search options in the form of drop-down boxes which allow them to opt to search by material type and collection for example. Post-search, there is an abundance of choice for limiting the search. The faceted navigation [3] functionality offers options to display just 'online resources' or just those that are currently 'available' (i.e. not checked out). Other options include refining by creator, topic or resource type. The number of likely results is given in parentheses next to each choice. The topic facets, based on subject headings from the library catalogue, can also be used as a starting point for completely new searches if required. If a work is available in more than one library

or in different collections, they can also be displayed as facets to choose from. Other facets are also possible depending on the fields originating from the MARC records in the catalogue, such as date of publication, language, classification, etc.

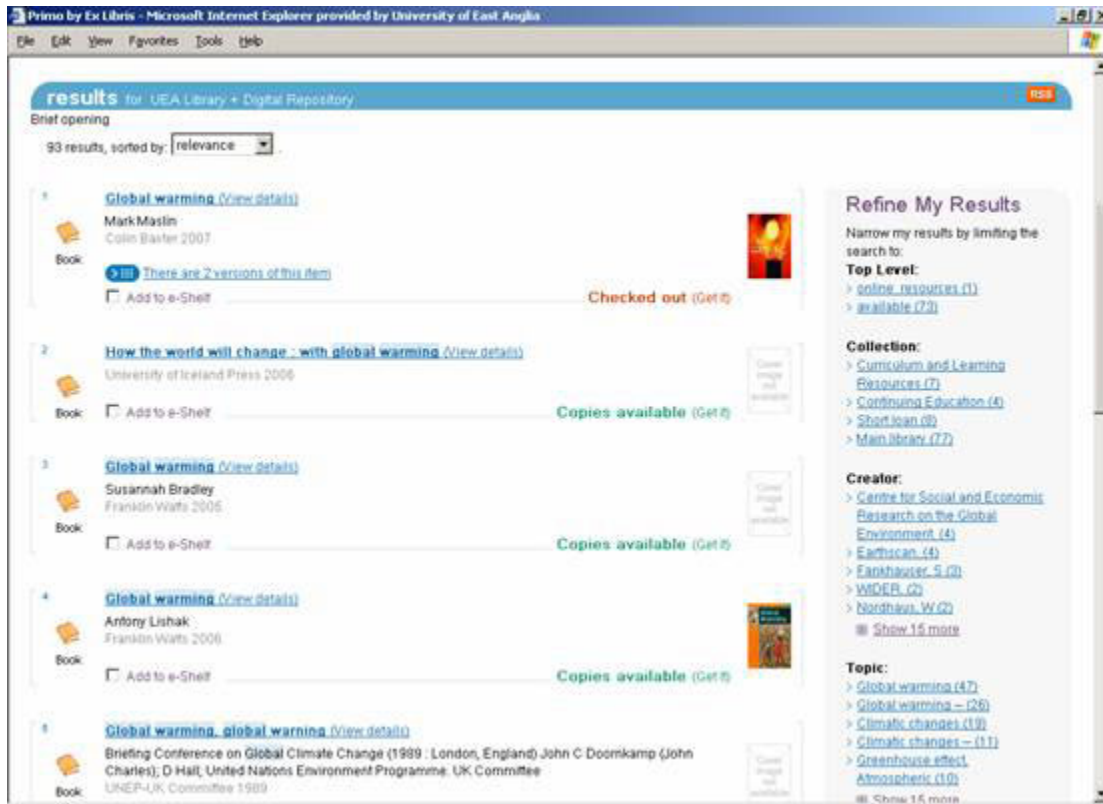


Figure 2: The PRIMO user interface: brief results with faceted navigation options to refine results

In order to provide a less cluttered approach to finding books, Primo includes functionality to group-related records under a single record. This so-called 'FRBRisation' [4] works well, particularly for textbooks where different editions are held under the one record. There is also a de-duplication option to bring together duplicate bibliographic records and merge them together which could be helpful for consortia set-ups.

From the brief results page users immediately receive an indication of whether copies or online versions are available. In version 1 of Primo the availability information is not real-time, being based on an availability pipe that runs at intervals throughout the day. Therefore in theory it is possible that if there is just one copy of a particular book, the holdings information could be out of date at that point. However, the 'Get it' link to the right of each record can be made to provide a quick link to what we have renamed 'Live holdings', which is the holdings information on our live library catalogue (Aleph) at UEA.

The 'Get it' function can also offer the user other appropriate copies or services. The 'Other services' tab uses your institution's link resolver service (e.g. SFX) to show links to other library catalogues, search engines or Google Scholar for example.

If users prefer, they can choose to expand a brief record to display the full record. This gives additional bibliographic information and also links to other services such as the table of contents from Amazon or an abstract from Syndetics Solutions [5] or the record in WorldCat [6]. This additional post-search enrichment of records, which supplements data from the original MARC records, is included as part of Primo, subject to any agreements needed with those third-party providers.

As you would expect from a major database, there are options to save, email and print results – with e-shelf (marked list) functionality to select records. Currently there are options to save to third-party applications such as del.icio.us [7], Connotea [8] and Refworks [9], with other export options anticipated in future versions. Other Web 2.0 and social networking features such as tagging and reviews (for logged-in users) are not particularly ground-breaking but nonetheless are what users now expect. If logged in, users can also save previous queries and set up alerts or RSS feeds for specific searches which can be configured only to update when new content becomes available.

The search functionality is more forgiving than a traditional library catalogue. There is ‘Did you mean?’ option which recognises common misspellings and offers options (e.g. if you type ‘mangment’ it will ask you if you meant ‘management’.) There is also a ‘still haven’t found what you are looking for?’ section which provides users additional guidance on what to do next if their initial search has been unsuccessful.

At UEA we have set up our default local search ‘scope’ to include our local library catalogues: the University of East Anglia and the Norfolk and Norwich Hospitals Library catalogues. We have also included a search of our local institutional repository (Ex Libris’s DigiTool). This is beneficial in terms of making our repository articles more visible as we look to promote the repository more widely. For ease of use, we also allow the option of searching the repository on its own by clicking on a separate tab which shows a new view just restricting the search and results to records in the repository.

The third tab is the remote search view, used for searching the library’s external databases and e-journals. The remote search option uses the federated search functionality from Ex Libris’s MetaLib. The user can search pre-selected groups of databases, or ‘quicksets’ as they are called, based on their subject matter or type of database. The databases are pre-selected by librarians and so the user does not know in advance what databases they are searching, just that the sets include databases relevant to their subject. There is a short delay whilst MetaLib pulls in and converts the records from these databases into its standard display format. If users get tired of waiting, there is a ‘show me what you have so far’ option. The users are then presented with a relevance-ranked, merged set of results from those databases.

In this version of Primo, it is not possible to include the remote search functionality within the same tabs as the local collections. There is a good reason for this because this remote searching relies on federated searching which is unlikely to return results as quickly as local databases. Therefore to combine local and remote searches might lead to unacceptable delays.

The remote search functionality works particularly well with bibliographic databases that index journal articles and book chapters. From the brief results, whose records look identical in format to the results from local searches, the user can quickly see whether or not the full text is available and, if it is, simply click the ‘Get it’ link. Primo then uses linkage software functionality (in our case SFX) to deliver the full text immediately to the user either within a frame of Primo, or, if the user prefers, in a new browser window.

If no full text is available from UEA, then the ‘Online resource’ tab is replaced by an ‘Additional Services’ tab which is available to offer other services such as a table of contents, Google or another library search. Primo also enables the user to choose to return more results from each database, thereby partially overcoming one of the main limitations of federated searching: the fact that only the first few records from each database are pulled in after an initial search.

Behind the Scenes

This section offers an overview of what it has been like to implement Primo from a project management perspective and the workload required to maintain it. Ex Libris provides a customer profile for Primo which outlines the various strands for setting up Primo. In basic terms these are: system administration and authentication, data management configuration, user interface configuration, and interoperability, i.e. integration with other systems. In our case those other systems are the SFX link resolver, MetaLib for federated searching and our digital repository, DigiTool. We are also planning to embed Primo search within our institutional portal. We realised from the outset that we would need a multi-skilled team for this project, drawn from across the Library and IT departments.

For ease of use, the majority of the configuration for Primo can be done through a Web-based interface called the Back Office. The configuration process for the user interface is fairly straightforward using these tools and Ex Libris provides

guidance in both the customer profile and Primo documentation. The only downside of the Web-based approach is that administrators do not get to see the underlying directory structure. This means it can be hard to assess how one aspect of the setup relates to another. The approach is good for those who like to start with the detail and work from there; others might prefer the overview first.

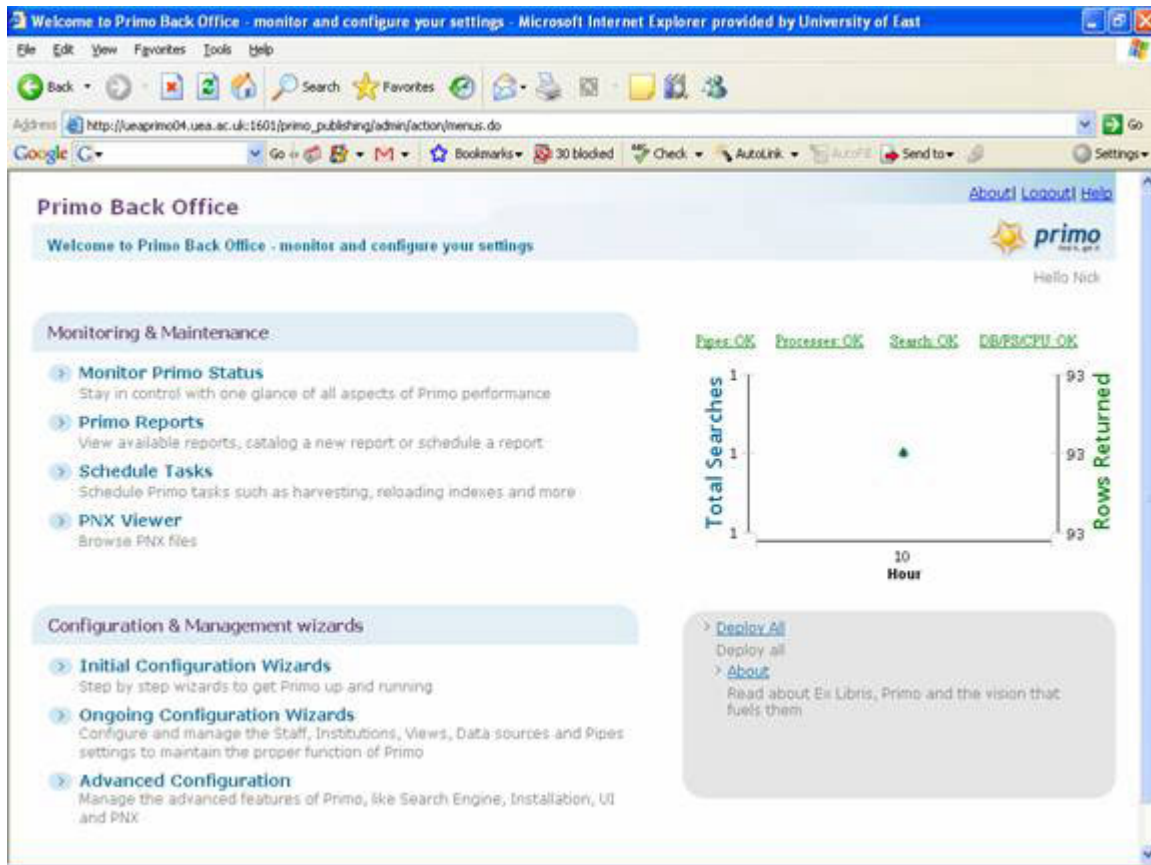


Figure 3: The PRIMO Back Office interface

At the heart of the Back Office are the tools for managing the Primo Publishing Platform which is where records from different sources are brought together. Although there are pre-existing pipes (transformations) for loading data, from your catalogue and repository, for example, it can be a complex process. Whilst the pipes work 'out of the box', the default settings are inevitably for a standard library with a standard setup. Most libraries have at least some local non-standard practices. Add to that the question of multiple branches or multi-site libraries, each with slightly different cataloguing practices, and some interesting obstacles arise. Having said that, Primo is designed to cope with consortium setups and has the flexibility to set up different views based on the user's context. This works well for libraries on completely different sites. However, challenges may remain for dealing with sub-libraries or branches if they share IP ranges or parts of your existing catalogue records.

So how do you refine the 'out-of-the-box' configurations? In order to display data from your catalogue, repository or other local databases, records need to be published into XML format and then each XML record needs to be 'normalised' to show in Primo. Normalisation creates an enhanced version of the XML record from each harvested source record. This enhanced record is called the PNX record (Primo Normalised XML). For this process you need someone who really understands the history of your library and its cataloguing quirks and foibles. In effect, you need to choose the person you would rely on if you were planning a library management system migration.

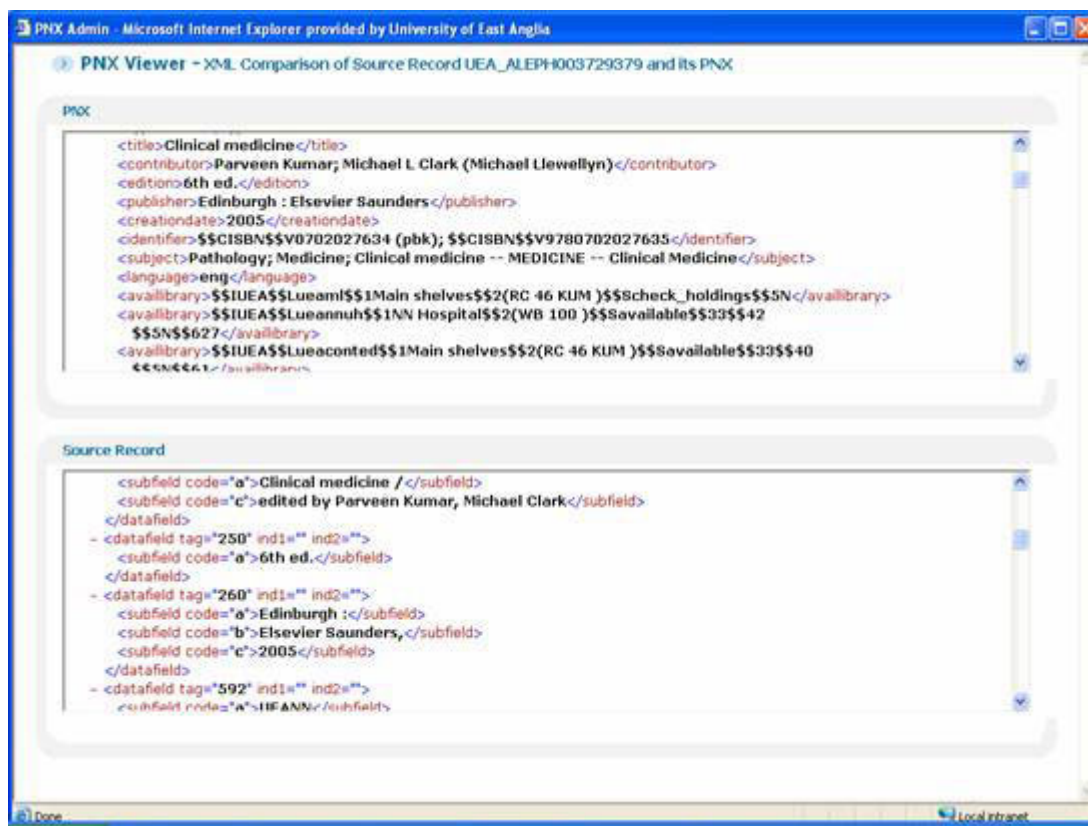


Figure 4: The PNX record viewer which shows the differences between a source XML record and a Primo PNX record

To make adjustments to a normalised PNX (Primo Normalised XML) record, you need to make changes to each standard pipe. You need to look at what fields are displaying in the Primo user interface, and how they are displaying, and make adjustments to their configuration accordingly. Since the Primo Publishing Platform has its own indexes, any changes are only reflected by completing a re-indexing process, called renormalisation, and this takes at least an hour to run depending on the size of your database. If you make changes to the data in your library catalogue to fit better with Primo, then you will need to reload, as well as re-index.

To save time, you can set up a range of test records which means you do not have to wait for renormalisation to see the effects of any changes you make. Some adjustments were straightforward, others have required the creation of rules to display the original data in different ways in the user interface, depending on the context. In some situations we also needed to change our current cataloguing practices in our library system so that the data worked well in both our existing catalogue and Primo. For example, we had to edit notes on our 856 records (our e-resource URLs) so they made sense in both systems.

As well as including data from the original MARC or XML record, the PNX record also has additional functional fields to enrich each Primo record with Web 2.0 and other value-added functionality such as facets and thumbnails. By working systematically through the data configuration we found we made good progress. This is because once a configuration issue is fixed, it is in theory fixed for ever (unless you change your cataloguing practices or library management system of course).

For the institutional repository integration into Primo, the preferred method of publishing is to turn on OAI (Open Access Initiative) publishing in your repository, in our case DigiTool. The DigiTool records are pulled using one of Primo's standard pipes for OAI harvesting and the records are then displayed in the consistent Primo format. As with the library catalogue, the standard pipe works well for this and adjustments can be made to the PNX records depending on the metadata format (MARC or Dublin Core) you have chosen in DigiTool. Depending on what is stored in your repository, Primo can present

images and other digital objects as well as links to digitised documents. The actual delivery of the objects, such as full text journal articles, is done by whatever system hosts that digital object.

All the local indexes on Primo need to be updated with new material on a regular basis, for example once a day. This process can be automated just to pull the new and changed records from each of the source databases.

Configuring the remote search functionality involves creating some additional Quicksets in MetaLib and then switching on various options in the PRIMO Back Office interface. The implementation of Ex Libris's proprietary Patron Directory Services (PDS) authentication (now available across all its products) is important for the remote search because this will determine access rights for the sets of databases to be searched. Depending on users' location or whether they have logged in or not, an indication of access rights will be shown to them for each of the quicksets prior to carrying out a search. PDS also enables single sign-on between all Ex Libris products and can link to external institutional authentication services using LDAP or Shibboleth for example.

Primo is not just a standalone product. It is interoperable with other library management products and software, not just Ex Libris products. In terms of where Primo appears in your institutional landscape of interfaces and services, it can easily be embedded in or deep-linked to. Initially we will be looking to make the search box available within our institutional portal and course management systems and at a later stage we will explore the potential of other Web services it provides.

In terms of system hardware, sizing obviously depends on the number of concurrent users and records. UEA is sized for 200 users and up to 5 million records. From May 2008, UEA will be running PRIMO version 2 on several HP BL460C Blade Servers each with 2 x 3GHz Quad-Core Xeon E5450 Processors. We are choosing to run on Linux rather than Solaris.

In terms of the total cost of ownership, Primo is a very flexible product and, as a result, it initially needs a high level of resource to maximise its potential. We expect this to reduce as Primo comes of age. Future difficulties may not lie so much in the maintenance of Primo itself as in its connections with the underlying source databases as they change and develop.

The Future

There are a number of other next-generation library search and retrieval products available, Encore [\[10\]](#) and Aquabrowser [\[11\]](#) being the most similar to Primo. Like Primo, they all aim to work with a variety of library management systems and federated and linkage services. One of the key challenges yet to be met by any of these products is the true integration of local and remote resources. Encore and Aquabrowser apparently show results on the same page as the local results but in a different frame; Primo shows the results in a different tab. We expect that, in future versions, users could choose to combine records from both local and remote searches, though perhaps at a later stage in their search, once more of the remote records have become available. In the longer term, the aim must be to obtain permission to harvest the actual journal data from third-party suppliers, such as Science Direct, which would speed up the search process and allow more accurate relevance ranking based on all the records. The barriers to this may be more political than technical.

Primo version 2 is due for release in the 2nd quarter of 2008. Looking at competing products, we hope to see more functionality such as an option for showing books by popularity, support for more material types, and more tools for integration and embedding Primo in other products, such as an alternative search screen for mobile devices. For smaller libraries, maybe a remotely hosted Primo could be offered for institutions that cannot afford the overhead of their own servers. We also expect to see other technologies like AJAX [\[12\]](#) being used, perhaps to cut down on the amount of data that needs to be loaded initially or to provide more real-time availability.

We hope to see additional support for dealing with local practices, so that some of the more subtle inter-relationships between records can be retained; for example if your library has analytic or hierarchical records. This may be particularly relevant to those who want to use Primo with less standard databases or with databases for which there are, as yet, no pipes available.

In its current version, whilst Primo aims to replace the search and retrieval interfaces of your local or remote databases, it does not claim to replace other functionality, such as the ability to place reservations. Should that be one of its longer-term aims?

Conclusion

Searching on Primo is closer to the experience of searching interfaces like Amazon than a standard library catalogue. Local and remote records appear in a consistent format which offers a continuity to the product as a whole which users may well prefer. The relevance ranking combined with an abundance of Web 2.0 options will be applauded by those who enjoy browsing as well as searching. It also provides valuable resource discovery benefits. Primo scores very highly in its scope, not only providing an alternative library catalogue, but also integrating e-resources, local articles and external resources. The quality of the source products such as SFX, MetaLib and DigiTool clearly helps to give Primo an edge, especially for existing Ex Libris customers.

In terms of other sites in the UK, at the time of writing, the University of Strathclyde and the British Library are the two other UK customers. However there are already pilot installations at several institutions internationally [13] and, to date, approximately 75 libraries and universities are signed up worldwide.

Ex Libris is clearly investing heavily in Primo and the architecture is likely to be used in the development of its next-generation library management system. We anticipate rapid development of Primo over the next few years. It is hard work being an early adopter but harder still would be confronting the likelihood that, without products such as Primo, our systems and services may be perceived as increasingly irrelevant.

References

1. OCLC. (2005) *Perceptions of Libraries and Information Resources*. Retrieved from <http://www.oclc.org/reports/2005perceptions.htm>
2. For example, use of Amazon Web Services <http://www.amazon.com/gp/browse.html?node=3435361> is free but requires a prominent link back to Amazon from your site.
3. For more information about Faceted Classification, see: Murray, P. (2007) *Faceted Classification of Information*. Retrieved from <http://www.kmconnection.com/DOC100100.htm>
4. For more information about FRBR (Functional Requirements for Bibliographic Records) see: UKOLN. (2005) *Bibliographic Management Factfile: FRBR (Functional Requirements for Bibliographic Records)*. Retrieved from <http://www.ukoln.ac.uk/bib-man/factfile/cataloguing-rules/frbr/>
- Danskin, A. & Chapman, A. (2003) Bibliographic records in the computer age. *Library & Information Update*, 2(9), 42-43. Retrieved from <http://www.cilip.org.uk/publications/updatemagazine/archive/archive2003/september/update0309d.htm>
- Riva, P. (2007) Introducing the Functional Requirements for Bibliographic Records and Related IFLA Developments. *Bulletin of the American Society for Information Science and Technology*, 33(6), 7-11. Retrieved from <http://www.asis.org/Bulletin/Aug-07/Riva.pdf>
5. Syndetic Solutions <http://www.syndetics.com>
6. WorldCat <http://www.oclc.org/worldcat/default.htm>
7. deli.icio.us <http://del.icio.us>
8. Connotea <http://www.connotea.org>
9. RefWorks <http://www.refworks.com>
10. Encore <http://www.iii.com/encore/>
For a sample implementation see Glasgow University <http://encore.lib.gla.ac.uk/iii/encore/app>
For in depth review, see Breeding, M. (2007) Encore. *Library Technology Reports*, 43(4), 23-27.
11. Aquabrowser: <http://www.aquabrowser.com/> For a sample implementation see Columbus Metropolitan Library: <http://catalog.columbuslibrary.org>
For in depth review, see Breeding, M. (2007) AquaBrowser. *Library Technology Reports*, 43(4), 15-18.
12. AJAX (Asynchronous JavaScript and XML).

13. Vanderbilt University <http://discoverlibrary.vanderbilt.edu>;
the Royal Library of Denmark http://primo-7.kb.dk/primo_library/libweb/action/search.do?vid=KGL;
the University of Minnesota http://prime2.oit.umn.edu:1701/primo_library/libweb/action/search.do?vid=TWINCITIES;
the University of Iowa http://smartsearch.uiowa.edu/primo_library/libweb/action/search.do?vid=uiowa

Author Details

Nicholas Lewis

Library Director
Information Services Directorate
University of East Anglia

Email: nicholas.lewis@uea.ac.uk

Web site: <http://www.uea.ac.uk/is>

[Return to top](#)

Article Title: "Implementing Ex Libris's Primo at the University of East Anglia "

Author: Nicholas Lewis

Publication Date: 30-April-2008 Publication: Ariadne Issue 55

Originating URL: <http://www.ariadne.ac.uk/issue55/lewis/>

[Copyright and citation information](#) File last modified: Tuesday, 20-May-2008 16:00:58 UTC

Main Contents	Section Menu	Email Ariadne	Search Ariadne
-------------------------------	------------------------------	-------------------------------	--------------------------------

Ariadne is published every three months by [UKOLN](#). UKOLN is funded by [MLA](#) the Museums, Libraries and Archives Council, the [Joint Information Systems Committee \(JISC\)](#) of the Higher Education Funding Councils, as well as by project funding from the JISC and the [European Union](#). UKOLN also receives support from the [University of Bath](#) where it is based. Material referred to on this page is [copyright Ariadne \(University of Bath\) and original authors](#).