

Dave Price

Chapter 11: Implementation of a new integrated library system

1 Background

The library community worldwide has experience of replacing an integrated library system and it usually follows a recognized process. The main interest in this study has been the radically different approach we have taken in planning the procurement and implementation.

For most university libraries, the most important single IT service is their Integrated Library System or ILS (also often referred to as a Library Management System or LMS, particularly in the UK). Comprising a catalogue, in our case, predominantly for printed books, but also for maps, music, video recordings and electronic materials, an ILS will also provide functional modules to support their management. Commonly there will be modules for cataloguing, acquisitions, circulation (i.e. lending) and serials management. There will also be an Online Public Access Catalogue, or OPAC, the interface through which readers can discover materials, reserve them, renew loans and manage their account details.

In the early 1980s, Oxford ran a pilot programme to automate four lending libraries. This was immediately followed by a major project to introduce automation to the central Bodleian Library and this service went live in 1988. From the outset, the strategy was to provide a union catalogue for Oxford, known locally as Oxford Libraries Information Service (OLIS). Membership rapidly grew and there are now over one hundred OLIS libraries, comprising the University's central, faculty and departmental libraries, most of the libraries of the colleges and permanent private halls, the libraries of several affiliated institutions and the Centre for Oxfordshire Studies (part of the local public library system).

The ILS then chosen was DOBIS/LIBIS. At the time it was marketed by IBM as a 'product offering' rather than a full product: IBM did not take on responsibility for software development, but left this to the user community to whom they made the source code freely available. This was to have a significant impact for Oxford as it meant that our local support staff could respond rapidly to enhancement requests from users. Over the years the Oxford libraries became accustomed to an environment where changes, sometimes quite idiosyncratic, could be made to the basic software, engendering a culture where software was modified to fit existing practice, rather than reviewing business processes in the light of software

capabilities. Such changes and enhancements were often referred to as ‘Oxford specials’.

Integrated library systems are highly complex in terms of software functionality and data structures. Inevitably they do need replacing as they become obsolete or because better functionality and performance is available. Upgrading to a new system is always a major challenge and is commonly undertaken every 7-10 years. Geac Advance was selected as a successor to DOBIS/LIBIS in 1995 and implementation completed in 1997. The software reproduced many of the Oxford specials functionality of the previous system. Over its lifetime, many more were developed, though none so extensive as the automated stack request system, developed as an addition to the circulation module to enable readers from anywhere in the world to order items from different bookstack locations for delivery in different reading rooms. This finally went into full production in 2002.

By 2003, we knew that development of Advance had ceased, and though the manufacturers promised to support it until the last customer, we knew that once again we must seek a replacement. As before, we undertook a major procurement exercise in 2004-05, this time a two-stage tendering under European Union regulations. Considerable effort from many library staff and committees went into drawing up the tender document, which had 1,786 requirements. Seven of the eight interested vendors submitted proposals and gave a one-day presentation to library staff, from which three were short listed. Each vendor gave demonstrations, including complex workflow scripts under the close scrutiny of librarians, covering most aspects of the requirements document, over a period of five full days.

Implementation began in 2005. There followed a period of intense activity on the part of both the vendor and librarians, which included training 700 staff in use of the new system. Despite overcoming numerous technical obstacles and developing some impressive features, in 2008 we decided to call a halt to the implementation.

Why did we take this decision? If library systems are complex, then they are doubly so in the Oxford context. The system must be capable of supporting an extended, federal structure – more than a hundred different libraries, largely independent with widely differing policies, and with members able to belong to more than one library, often with different statuses. Very few systems on the market approach this level of complexity. On top of this were the aforementioned Oxford specials. The primary reason for the failure to progress to implementation was the excessive requirements demanded by Oxford. It had been unrealistic for us to expect a vendor to achieve in a short space of time the extensive development required to reproduce – and exceed – the functionality we had developed over two decades and two library systems.

2 Moving on: Resource discovery

In one sense, we were back to square one regarding the ILS implementation. However, it gave us the opportunity to focus on resource discovery. We had planned independently to introduce in 2008 two ancillary library systems: firstly, to replace our Open URL resolver with Ex Libris's SFX and, secondly, to replace our home-grown interface to electronic databases and allied resources with Ex Libris's MetaLib. Meanwhile, a new breed of interface, referred to generically as resource discovery platforms, was beginning to emerge. They differ from the traditional library search interfaces, OPACs, in several respects. The interfaces are more modern and are more like the interfaces our users encounter elsewhere on the internet. Running independently of the ILS and using dedicated hardware and modern indexing technology, typically Apache's Lucene/Solr, they are usually much faster than OPACs. Rather than traditional library searching methods, e.g. index browsing and precision (Boolean) searching, the new interfaces encourage users to start with a simple 'one-box' search, like Google, which may generate very many hits. These may then be refined by the selection of dynamically produced facets and scopes to assist the user in discovering what they need. This modern approach is augmented by relevance ranking, which attempts to present results most usefully to the user. So-called web 2.0 features provide the user with the ability to tag, review or 'favorite' records, or share with other services such as reference management software and social media sites. Importantly for Oxford, where many of our collections are not included on our ILS but are listed in other catalogues and finding aids, resource discovery platforms can take catalogue records not just from the ILS but from different sources and allow users to search across them together.

We determined that we could, without waiting for a new ILS, bring significant improvement to user experience by introducing rapidly a resource discovery platform. We chose Ex Libris's Primo, badged locally as 'SOLO: Search Oxford Libraries Online'.

We had in the past provided our users with highly customized software in terms of appearance and functions. But we had learned that this came at the price of high development costs and difficulties applying upgrades to software which had diverged significantly from the mainstream product. So to install Primo rapidly and to avoid software configuration problems, we chose to minimize customization and to use it as closely as possible to 'out of the bag', or the 'vanilla' system. Behind this reasoning was a desire to send a message to staff that we have moved on from a time when software can be highly customized and perfected before entering production. We were now in a 'beta-culture' age when software can be released before it has been fully refined and tested to exhaustion,

as could be seen elsewhere on the Internet. Users nowadays understand that this is the price of an early release – but also that in the new, agile environment, small incremental improvements can be made at short notice, without inconveniencing them or having to wait, for example, until each Long Vacation before changes are introduced.

SOLO was launched 1 October 2008 by which time both ‘OxLIP+’ (MetaLib) and ‘Oxford University e-Journals’ (SFX) had also been introduced. SOLO searches and displays not only records from our integrated library system (which includes e-journals and e-books), but also from our Chinese and Japanese catalogues, our electronic databases, and Refugee Studies database of grey literature.

3 Stepping back: External consultation

Though providing some breathing space, SOLO did not remove the need for a replacement ILS. To help us consider the next steps, we invited a colleague from Harvard Library to work as a consultant with us, sharing knowledge of and expertise in the ILS world to provide a neutral view of the Oxford experience and to suggest how we might progress procurement of a new system.

Agreeing that a replacement system was essential and urgent, he advised an ‘aggressive but not unrealistically tight timeline’ – enough time to ensure success but not unnecessarily long, since otherwise the project had the potential to consume as much time and resource as were available. This would require a highly concentrated effort for system selection and contract negotiation and good project management during the post-contract implementation phase. He warned against lack of precision in communications with the vendor regarding system configuration and database conversion and urged us to canvass other sites about their implementation experience, factoring that into our timeline. He also advised that responsibility for managing the vendor relationship should be at a high management level.

The consultant also warned us against ‘prolonged decision-making around implementation decisions’, noting that Oxford libraries have a strong tradition of consultation and consensus-building, in part reflecting the decentralized nature of the library system. This hit home as several hundred librarians had had significant involvement in the last procurement and implementation.

He made some acute observations regarding procurement. ‘The marketplace for ILS systems is smaller and weaker’. Whereas in the past there had been eight or nine potential suppliers, as a result of much consolidation and slow sales, there were fewer systems, perhaps just two, that could realistically provide for a large,

academic research library like Oxford. He noted, ‘the role of the ILS has shrunk’. This was not to decry its critical importance, but the ILS was not longer the single, central system in the library’s environment, but one component, albeit a major one, in an extending digital infrastructure, with many new initiatives requiring the library’s attention. Therefore the resources dedicated to the project should be considerably reduced. His view was that customization - the Oxford specials - was to be avoided whenever possible and instead we should put effort into planning how to adapt our workflows and processes in order to exploit better the software as delivered.

Notwithstanding this last point, he recognized that one area of functionality, our Automated Stack Request system, would not be provided by any of the systems on the market and would have to be developed. He advised that we undertake an analysis of the service requirements of book delivery in light of the forthcoming Book Storage Facility, its internal inventory system (BSFIS) and the interface between them, in order to provide vendors with precise specifications.

The consultant advised a faster, simpler and less resource-intensive selection process. Rather than the traditional exhaustive comparison of systems using a long list of functional requirements, he suggested making the decision using two elements: firstly, by appropriate site visits and by assessment of how successful an implementation has been at peer libraries; secondly, by shortlisting high-level, truly critical factors relating to the vendor and the system.

To this end, he met with the Bodleian Libraries’ Cabinet of senior managers and together they identified the following high level criteria:

- Scalability
- Configurability
- Integration with Resource Discovery
- Integration with Other University Systems
- Support for Automated Stack Request
- Implementation
- Maintenance & Support
- Supplier Relationship
- Supplier Viability
- Price

Another outcome was the establishment of a new post of IT Programmes Manager to manage both the ILS and the BSFIS projects. With great foresight, this was made a permanent post, recognizing that once these projects had been completed, there would certainly be more major IT developments waiting in the wings.

4 Marching Forward: the Procurement Process

Responsibility for the ILS and BSFIS implementations lay with the Integrated Library Systems Steering Group (note the plural form), itself reporting to the Bodleian Libraries' Executive. This group had earlier overseen the implementations of SOLO, OxLIP+ and Oxford University e-Journals. Led by the Associate Director, then heading up Special Collections and IT, it included the Assistant Directors of Research & Learning Services and of Collections & Resource Description, the Head of Systems & e-Research Service, and the chair of the Committee of College Librarians. The ILS Steering Group took forward the consultant's recommendations and in line with his advice set up a small, specialist group, the ILS Scoping Group, to run an accelerated and streamlined procurement, reporting monthly to the Steering Group.

The Scoping Group was heavily weighted with members of the Systems & e-Research Service, who would be effecting the implementation, but it also included representatives of Research & Learning Services, Collections & Resource Description and the college libraries. It first met in February 2009 and its brief was to deliver recommendations to the Steering Group in July.

There were several interesting aspects about this procurement which made it very different from the traditional approach we had employed in the past.

University regulations allowed us to an extent to define our own procurement methodology. Rather than publishing an open invitation to tender, provided we satisfied the University's procurement policies, we could restrict it to pre-selected vendors. As noted above, changes in the ILS marketplace meant that realistically, for our purposes, there were perhaps two leading contenders, so these went on our shortlist. We added to it the supplier of our current system because

- they were extremely familiar with our current practices;
- they had hitherto provided excellent service;
- their replacement product looked modern with great potential;
- we could be sure they could make a good job of migrating our data from the old to the new system;
- they had written our original Automated Stack Request system.

Our previous tender document had been 190 pages long with 1,786 requirements to which the vendors were expected to respond individually, specifying the extent to which and how they could meet them. Key to our methodology was to turn this on its head and to proceed from the premise that all three short-listed systems could meet our needs. Rather than our former approach, we chose to run the procurement as a 'competitive dialogue' where we could work together transparently with the vendors to ensure a common understanding of issues and help them produce their final submission.

Although Oxford had in the past been seen as a large, influential and prestigious customer that vendors would like to have on their books, it was generally understood that taking on Oxford would be challenging in terms of its complex federal structure and high level of customization. A major risk we identified was that the short-listed vendors would view Oxford as too challenging and might choose not to tender.

Early discussions with the three vendors in January 2009 outlined our intended procurement procedure and emphasised Oxford's desire to implement rapidly, to work as far as possible with standard software and to minimize customization and development. This was followed by a formal letter to the head of each company, which described the proposed process and timetable and asked if they were willing to participate. To our relief, they all were and in April they were issued with a *Competitive Dialogue Briefing* document just 35 pages long.

The first section provided background and objectives of the project and the evaluation and selection process. The second section provided a high-level description of the functional scope of the main system modules used by Oxford, noting that consideration should be given to Oxford's federal structure, the main areas being OPAC, resource discovery, cataloguing and authority control, acquisitions, serials management and circulation. Specific reference was also made to the Automated Stack Request requirement, now stripped to bare essentials, expecting that most vendors would need to do at least some development to enable their systems to provide this.

The third section took the ten key selection criteria referred to earlier. These were then broken down into 'evaluation points' and vendors were asked to address all forty. So, for example, the first, scalability, was broken down into system performance, search functionality, federalization and ability to accommodate structural change.

Having received the document, the vendors had one month to provide their draft responses. We invited them to contact us in that period whenever they wanted to seek any clarifications needed to help prepare their draft responses. All did, and the interchanges were fruitful on both sides. In addition, regular news bulletins were issued to them, so they each received any new information elicited from us by the others.

From the viewpoint of the vendors, the next stage was two days with the Scoping Group to scrutinize the response with the express purpose of improvement before its final submission. This included a meeting with the ILS Steering Group and the Executive concentrating on strategic issues.

Meanwhile, the Scoping Group conducted the equally important activity of customer site visits. For each vendor system, two sites were visited. The sites were provided with copies of the *Competitive Dialogue Briefing*, so they would under-

stand the nature of the procurement and could relate it to their own implementations and experiences. Each visit lasted two days. We should record here our immense gratitude to all of the sites who unreservedly gave their time and also unstintingly, their hospitality. Not only were the visits highly instructive, they were highly enjoyable as well and the professional contacts made have endured.

The Scoping Group took a holistic approach to the evaluation process, with the vendors' written responses considered alongside the knowledge gained from the reference site visits and vendor meetings.

Sections two and three of the *Competitive Dialogue Briefing* were expressly designed to provide a framework for the analysis of the suitability of the candidate systems. Section two sought to test the assumption that any of the candidate systems should be broadly capable of meeting our requirements and to identify any discrepancies and gaps in this functionality. Section three took the ten key selection criteria and expanded them into a series of evaluation points against which each of the three candidate systems could be measured. The numerical scores were weighted to reflect the relative importance of criteria and evaluation points. It is feasible that any of the three vendors could have provided for us – indeed that was the premise at the outset – but the outcome was the recommendation from the Scoping Group that Oxford move forward with the acquisition of Aleph from Ex Libris as a replacement ILS.

Though the Libraries accepted this recommendation in July 2009, we were not able to announce the result and commence contract negotiation with Ex Libris until we had formally secured funding from the University. This was awarded in September 2009 after the consideration of the business case. (Its preparation was the first task of our new IT Programmes Manager.) The delay this introduced was difficult for all the vendors and we were grateful for their good-humoured forbearance during this period.

One nugget from several of the earlier site visits had been to allow plenty of time for contract negotiation, and this indeed proved to be the case. Once we had informed Ex Libris of our decision (and shown due diligence by following University guidelines to run a procurement exercise to appoint an external solicitor), this took a full three months. This threatened our first milestone, the project kick-off in January 2010, but happily contract signing was completed in the nick of time and good will on both sides meant we were able to work with Ex Libris during the pre-implementation phase on project planning.

5 Project planning

Great care went into the overall implementation strategy. Obviously, the vendor hoped for as rapid an implementation as possible, and for a site like Oxford, they considered a year should be sufficient. Library management too were keen to have the new system running as soon as possible. On the other hand, library staff had toiled for three years during the previous implementation: they knew the scale of the detailed work that had to be accomplished and many felt three years might not be enough. But against this, it was evident that many were fatigued and it was clearly undesirable that they should suffer so soon another protracted implementation. Another factor was maintenance for the old system, which was only assured until July 2011. Accordingly a phased approach was devised.

In broad terms, the first phase would last eighteen months, from January 2010 to July 2011. At the outset we would have sufficient hardware in place for Ex Libris to commence their software installation and then to work with us configuring the system, training our support staff and running test data migrations from the old ILS to the new. By January 2011, they would hand over to us a completed system, as it would be when we would go live. In effect, this was the main implementation project as far as the vendor was concerned and would last one year. It is generally best for us to introduce new major systems over the summer, after exams and before the next academic year, to minimize impact on our users, so moving the new ILS into production was planned for July 2011. This would give us a period of six months with the new ILS effectively ready to go, when final data migration issues could be resolved, staff trained, and any outstanding glitches and structural and workflow issues could be addressed. It could also serve as contingency so implementation overruns could be absorbed by this ‘buffer period’.

What we expected to have at the end of the first phase was a system pretty much close to vanilla, but configured to provide for Oxford’s federal structure and to have requisite functionality to allow readers to call books from our stacks, interfacing with the BSFIS. We planned and resourced a second phase, lasting 18 months, to build up the system to its full functionality and for further ancillary development. Reporting provides a good example of the type of work allocated to phase two where over the sixteen-year lifetime of the old ILS a veritable plethora of specialized reports had been developed for the different libraries. We could provide a limited set of standard reports on day one. In phase two we are devoting development and training effort to introducing a sub-system, the Aleph Reporting Center (ARC), which will enable librarians to create many reports themselves to attain the levels of reporting they have enjoyed in the past. It will also still be necessary for systems staff to create bespoke reports to attain the full richness previously enjoyed.

6 Implementation

The existing ILS support team comprised three applications programmers, a stock control programmer and two library systems support staff specializing in training and configuration of the circulation and acquisitions functions, including serials registration. For the duration of the two phase implementation and also to support the introduction of the BSFIS (see Chapter 10), we established five additional three-year posts: an application developer, an Oracle database administrator, a data administrator, a project administrative assistant and a systems analyst providing support and training in cataloguing. All staff worked closely together as a single ILS implementation team organized on a task basis.

Phase one of the implementation proceeded as planned. After an initial installation of a test environment, in April and May 2010 Ex Libris trained our systems staff and helped them configure the software for Oxford. The full hardware was installed and the production environment built. Data migration from the old system to the new was a critical aspect of the implementation and, starting in June, three test migrations were performed in the first year, each run improving both in terms of data integrity and speed. We were able to sign an interim acceptance of the Aleph software in December 2010 on schedule, so from the viewpoint of Ex Libris they had completed the implementation project.

There was still considerable work ahead for our implementation team. The principal administrative units in the Aleph structure and, for Oxford, equivalent to independent libraries, known as ADMs, needed configuration. A final test migration was performed in April 2011. We felt this was necessary because the high volume of new data created by the barcoding and book move projects meant that our database had changed considerably during 2011. Staff training began in April 2011 and continued beyond go live in July.

Overall, the phase one implementation went according to plan and extremely well, and this success was due largely to the IT Programmes Manager's excellent project management, a good relationship with Ex Libris and their project team, the skill and dedication of our implementation team, a tight governance structure and the support of senior management. Some decisions and aspects are worth highlighting for their intrinsic interest and where lessons were learned.

As regards training, a strategic decision was that, rather than hiring trainers specifically for the project, we should use our own staff (so retaining knowledge) and provide less broad-brush training for staff in general, but more focused training for staff regularly using the system. Also that the trainers would produce good documentation and strong follow-up support once the new system had gone live.

Another important strategic decision was to treat the management unit, the Bodleian Libraries, as a single structural unit (ADM in Aleph terminology) rather

than a large number of individual units as hitherto. In comparison with the other administrative units, this has added complexity, but it has added significantly to ease of administration in terms of support for centralized functions, circulation, reports and budgetary control. It has also provided a structure which would allow the move to harmonize policies across the Bodleian Libraries in terms of readers' statuses and privileges if this is considered desirable.

A significant step forward from the viewpoint of Oxford users with the introduction of SOLO and Aleph was the adoption of Oxford single-sign-on for access to library-provided electronic resources and patron accounts. This meant that University members are now able to use the same authentication for our services as they do for the University's email, virtual learning environment and many other services. This improvement has been extremely well received. However, Oxford SSO is available only to University members and a high proportion of Bodleian Libraries' membership is non-University. The corollary was that we had to develop and support a parallel authentication system for 'Library Card' holders (non-University members), which was a significant task.

We chose to virtualize our hardware architecture, basing it on VMware ESX, in order to provide a highly available, resilient, dual-site infrastructure. This added topological complexity but Ex Libris were able to configure the Aleph software successfully to take advantage of it. Behind the Aleph software lies an Oracle database. Initially, we positioned the two mirrored Oracle database servers within our VMware infrastructure. During the implementation we experienced problems but discovered that Oracle have not certified their software to run in a virtualized environment. To mitigate this risk in the future we purchased and installed dedicated Oracle servers outside VMware.

By replacing our old ILS with Aleph, using the OPAC via Primo (OvP) technology we could in principle provide our resource discovery service, SOLO, with full patron functionality in terms of the ability to place holds and reservations, renew loans and check account details. Beforehand, with SOLO acting as a front end to Aleph's predecessor, a user would have to 'shell out' to the Geac Advance OPAC to do these. This raised the question of whether we should now move ahead with SOLO as the single interface to the ILS or to make the Aleph OPAC available as well. It was important to make the decision quite early in the implementation. Advantages of a single user interface included ease of navigation for users, reduced systems support and reduced user support and user education. Against this was the current lack of index browsing and date-range searching, and the difficulty experienced by a section of our readers in searching and retrieving because of the way search results are clustered.

The ILS Steering Group, which comprised high-level representatives covering the Oxford library community, considered carefully the pros and cons, knowing

that the final decision would not suit everyone but that it was in everyone's interest to make it rapidly. Their recommendation to Executive was to go live without a dedicated OPAC. There has been some dissatisfaction, chiefly from librarians. The SOLO User Group, made up mostly of reader services staff but including systems support staff, has achieved much by analysing SOLO issues, providing development suggestions both for our development staff and Ex Libris, and importantly providing and promoting new search strategies to tackle perceived deficits, which has been extremely helpful, particularly to those used to traditional OPACs and coming to SOLO for the first time.

The ILS implementation was not the only project taking place at this time. The barcoding and bookmoving projects in particular required significant amounts of work to be done by the ILS implementation team. Here, alignment and co-ordination of potentially competing or conflicting project plans were dealt with most effectively by the Bod Squad, attended by the IT Programmes Manager (see chapter 2).

Inevitably there has to be a period switching to the new ILS when data is frozen. Working closely with Ex Libris and refining data migration procedures, we were able to bring this down from an original estimate of three weeks to eleven days – a working week and the two weekends either side. We scheduled shutting down the old ILS at 5 pm 7 July 2011 and then exporting its data. During that very day, the old ILS server failed and we were unable to restart it. In fact it took five days to recover the old system sufficiently for us to begin the data export. In all it added four days to the migration period and we finally brought Aleph into full production 22 July 2011. At the time it felt devastating but in the context of a major project such as this, four days' slippage is not so bad. A major business system such as an ILS should not be down for five days. At least this incident supported our view that we needed a replacement ILS!

We had chosen to go live at the beginning of the summer rather than the end in order to deal with any teething problems before the start of full term in the coming academic year. In fact, this bedding in period was busier than we had expected. The sophisticated automated stack request system we used with Geac Advance had now been replaced with the Aleph closed stack management sub-system which had been enhanced by Ex Libris to match our requirements. After go live, we had to make many configuration adjustments, especially to loan limits, to fine-tune it further for our complex environment.

Another area of concern was the ability of the system to handle notes describing particular copies of books within the bibliographic record, which are of particular importance in Oxford to those working on early printed books. Following some modifications by Ex Libris, the implementation team was able to devise a suitable protocol which was released over the summer.

Having dispensed with an OPAC, SOLO had to deal with much more traffic and transactions, with a consequence that we experienced performance problems.

We dealt with this by installing more memory. Unlike Aleph, which is a mature product, Primo, the underlying software of SOLO, was relatively new at that time and under rapid development. Inevitably we encountered some bugs and constraints, particularly in delivering OPAC function via SOLO, that needed immediate attention, investigation and fixing.

That we should have such problems was not a surprise. Where perhaps we should have invested more thought and effort was in expectation management of the Oxford librarians for the immediate post-go-live period. There was amongst many a sense of disappointment that the new system was not as good as the old. This was exacerbated by the four-day delay in going live, the teething problems mentioned above and for some, regret at losing a traditional OPAC. In retrospect we think that we had not succeeded in imparting several key messages. Firstly, that to avoid a prolonged implementation, we had broken it into two phases, the first delivering a basic but nevertheless a full production service. There would be a resourced second phase to further build and extend the system. There were mechanisms in place for library staff to influence the priorities, and report generation would certainly be one of them. Secondly, we can expect teething problems when a major business system is brought into production in a complex environment such as the Oxford libraries. The project plan allowed for three months over the summer, which for most libraries is a quiet period, to deal with them. Thirdly, we do believe that a modern resource discovery system that can incorporate different catalogues and finding aids, not just the ILS, in place of a traditional OPAC, is the right strategy. Systems like SOLO are under rapid development and their apparent deficits are being made good, or can in many cases be obviated by an open and imaginative approach. Fourthly, highly customized software is no longer sustainable. Rather than reproducing our traditional workflows, we should where necessary be examining them and redesigning them in the light of what is built into the software.

Phase two began in earnest in November 2011, once the teething period was properly behind us, and it is currently resourced for a further 18 months. The ILS Steering Group, representing all areas of the Oxford library community, manages the implementation team's priorities. A number of them have already been accomplished: a library control system which will allow a library itself to make configuration changes, such as opening hours, loan periods; 'Scan and Deliver', a new service to enable University members to request scans from items held in the BSF; and setting up a lending collection within the Gladstone Link in the central Bodleian Library. Other developments under way include a SOLO interface for mobile devices; roll-out of the ARC reporting system to the libraries;

using metadata from other catalogues and finding aids to search and retrieve in SOLO; further extension of reporting and record exports; self-return; the addition of several new member libraries; migration of the Chinese and Japanese catalogues into Aleph; and disaster recovery planning and testing.