

Go with the flow: data management and synchronisation across systems at the State Library of Victoria

Simon Huggard
Applications & Digital Services Manager
State Library of Victoria
shuggard@slv.vic.gov.au

Michele Hosking
Library Applications Administrator
State Library of Victoria
mhosking@slv.vic.gov.au

Abstract:

From 2008-2010, the State Library of Victoria has implemented five enterprise-wide library systems as part of the SLV21 project planning framework. A considerable amount of time and effort has been put into working with data flow across these different systems to ensure that staff workflows and data use and re-use is properly managed. This paper will cover a range of data modelling and planning that has been necessary to work with each system. Some best practice methods are discussed and outcomes and recommendations put forward.

Introduction

The State Library of Victoria has implemented five major enterprise library systems between January 2008 and February 2010. These systems were funded in part by a special grant from the State Government in 2005/6 which established the Office of E-Strategy and the SLV21 set of projects, many of which have transformed library processes, services, planning and strategies. This paper focuses on some of the technology/systems aspects of these projects.

Four of these systems (Primo, DigiTool, Metalib & SFX) are systems provided by the vendor Ex Libris and the other is the Relais Interlibrary Loans System from Relais International. There is another enterprise-wide system within this mix: Voyager, also from Ex Libris, which the library implemented in 2001. Other smaller (not necessarily enterprise-wide) systems are also part of the puzzle, such as EAD XML documents and the open source system Archivists Toolkit, both of which have been implemented within the same time frame.

Each of these systems uses the same or very similar bibliographic, holdings, patron and other data from single or multiple sources, which the library needs to manage in some way, in order for there to be:

- consistency across the user experience (particularly for authentication)
- consistent descriptive (bibliographic and holdings) metadata across systems
- accurate requesting and ordering across each system
- consistent and accurate searching and display of data within each system
- accurate printing and referencing and identification of objects coming out of each system
- accurate referencing of objects for external harvesters/crawlers/data repositories

This paper will cover a range of data modelling and planning that has been necessary to work with each system. We will also outline the workflows, across and within each system, that we have configured or developed in order to achieve the goals outlined above. Because the State Library of Victoria is implementing a diverse range of library service systems all at the one time, that are complex and are interrelated, it is imperative that we get the data flow working between these systems for our services to continue to operate successfully.

As a cultural institution with a diverse history and mission to provide services to all Victorians, the State Library of Victoria has some interesting challenges in terms of service provision and must therefore implement specialised and customised products to meet the needs of our patrons.

Some best practice methods are discussed and outcomes and recommendations put forward.

Current research/publications

The challenges of managing descriptive metadata and synchronisation of data have been concerning the library community for some time. This is particularly true in relation to descriptive metadata and holdings data in the MARC format. More recent challenges have arisen, however, due to the expanded nature of work that libraries are now involved in, and the challenges of working with the different systems required to deal with different work roles and responsibilities.

Apart from traditional library catalogues, there are many examples of new systems which library (and non-library) staff have recently had to deal with in their expanded work portfolios. Examples in the academic community include: learning management systems, learning object repositories, educational/teaching management systems, information portals, student and staff registration systems, raw research data management systems, repositories to record and manage (and account for) research output, subject and discipline related data archives and image repositories. In the area of cultural institutions, examples include: heritage accessioning and management systems, image repositories, conservation databases, exhibition management systems, and customer relations management systems.

With the recent emergence of add-on and stand alone systems to support a desired Web 2.0 functionality, there are also endless possibilities of further systems being involved in the mix of data use, re-use, mash ups, alerts, personalisation and so on. This paper will not attempt to catalogue the myriad of systems and requirements in these cases, except to flag that this is another issue that will need consideration into the future.

The proliferation of these systems makes it even more important to have the correct systems and information architecture in place to deal with data management and synchronisation. Current research shows that library professionals are very much aware of the challenges raised by this abundance of new systems and the need to manage such systems effectively. The key challenges include: ensuring data consistency between systems, minimising effort in terms of data synchronisation, mapping of data elements across systems, maintaining the integrity of data semantics, and ensuring that the repurposing of data is clearly understood and implemented successfully across diverse systems. A simple example of the latter might be a creator/author in a digital repository versus an authorised personal name in a library catalogue.

Much has been written on the need to re-use and repurpose the data held in libraries' catalogues. Calhoun's challenge to traditional bibliographic control practices and to traditional methods of providing access to bibliographic data (Calhoun 2006) served as the catalyst for debate about the worth and future of traditional catalogues. The bibliography at the end of this paper lists a variety of contributions to that discussion, from Mann's passionate response to Calhoun (Mann, 2008), to Christen and Tennant's argument for integration of disparate resources.

Christensen and Tennant identified integration of access as the primary issue, asserting that:

Integrating access to disparate information resources is a continuum ranging from completely integrated to geographically dispersed and technically divergent. All things being equal, the best situation for the end user is a completely integrated system, with all appropriate metadata stored internally in a common format uniformly applied. This provides the greatest flexibility to meet the needs of a specific user group or purpose, since total control (both just-in-case and just-in-time) can be exerted on the system. (Christensen and Tennant, 2005).

Christensen and Tennant considered complete integration to be the only acceptable option:

All solutions other than complete integration are a compromise, in which the control you have over all aspects of the system may be considerably lessened, from user search options to how results can be manipulated and displayed. (Christensen and Tennant, 2005).

In order to achieve complete integration, they identified a number of “integration principles”:

- Metadata should only be created, not destroyed
 - Integrate metadata whenever possible
 - Exploit metadata similarities
 - Honor metadata differences
- (Christensen and Tennant, 2005, p.5-6).

This concern was shared by the University of California Bibliographic Services Task Force, who agreed that:

We need to look seriously at opportunities to centralize and/or better coordinate services and data, while maintaining appropriate local control, as a way of reducing effort and complexity and of redirecting resources to focus on improving the user experience. (University of California Bibliographic Task Force, 2005, p.2).

In her report to the Library of Congress on the future of bibliographic control, Boehr said that redundancy of cataloguing effort should be reduced by repurposing existing metadata (Boehr, 2008). We have applied this principle and recognised its role in also reducing data redundancy.

In 2004, Andrew K. Pace noted that very few of the value-added and discovery systems “completely integrate with the legacy system. Libraries are forced to take these standalone products they have created or bought and hack access to the main system through use of APIs, clever Perl scripting, and scheduled server jobs that only mimic true interoperability” (Pace, 2004). Despite significant advances in both systems and interoperability since that time, Pace’s comment is still appropriate, particularly in our case at the State Library of Victoria, due to the diverse nature of our collections and the corresponding complexity of our data and services.

However, it appears that little has been published on data management and synchronisation in the context of libraries, library data and increasingly diverse library service delivery systems.

Data modelling

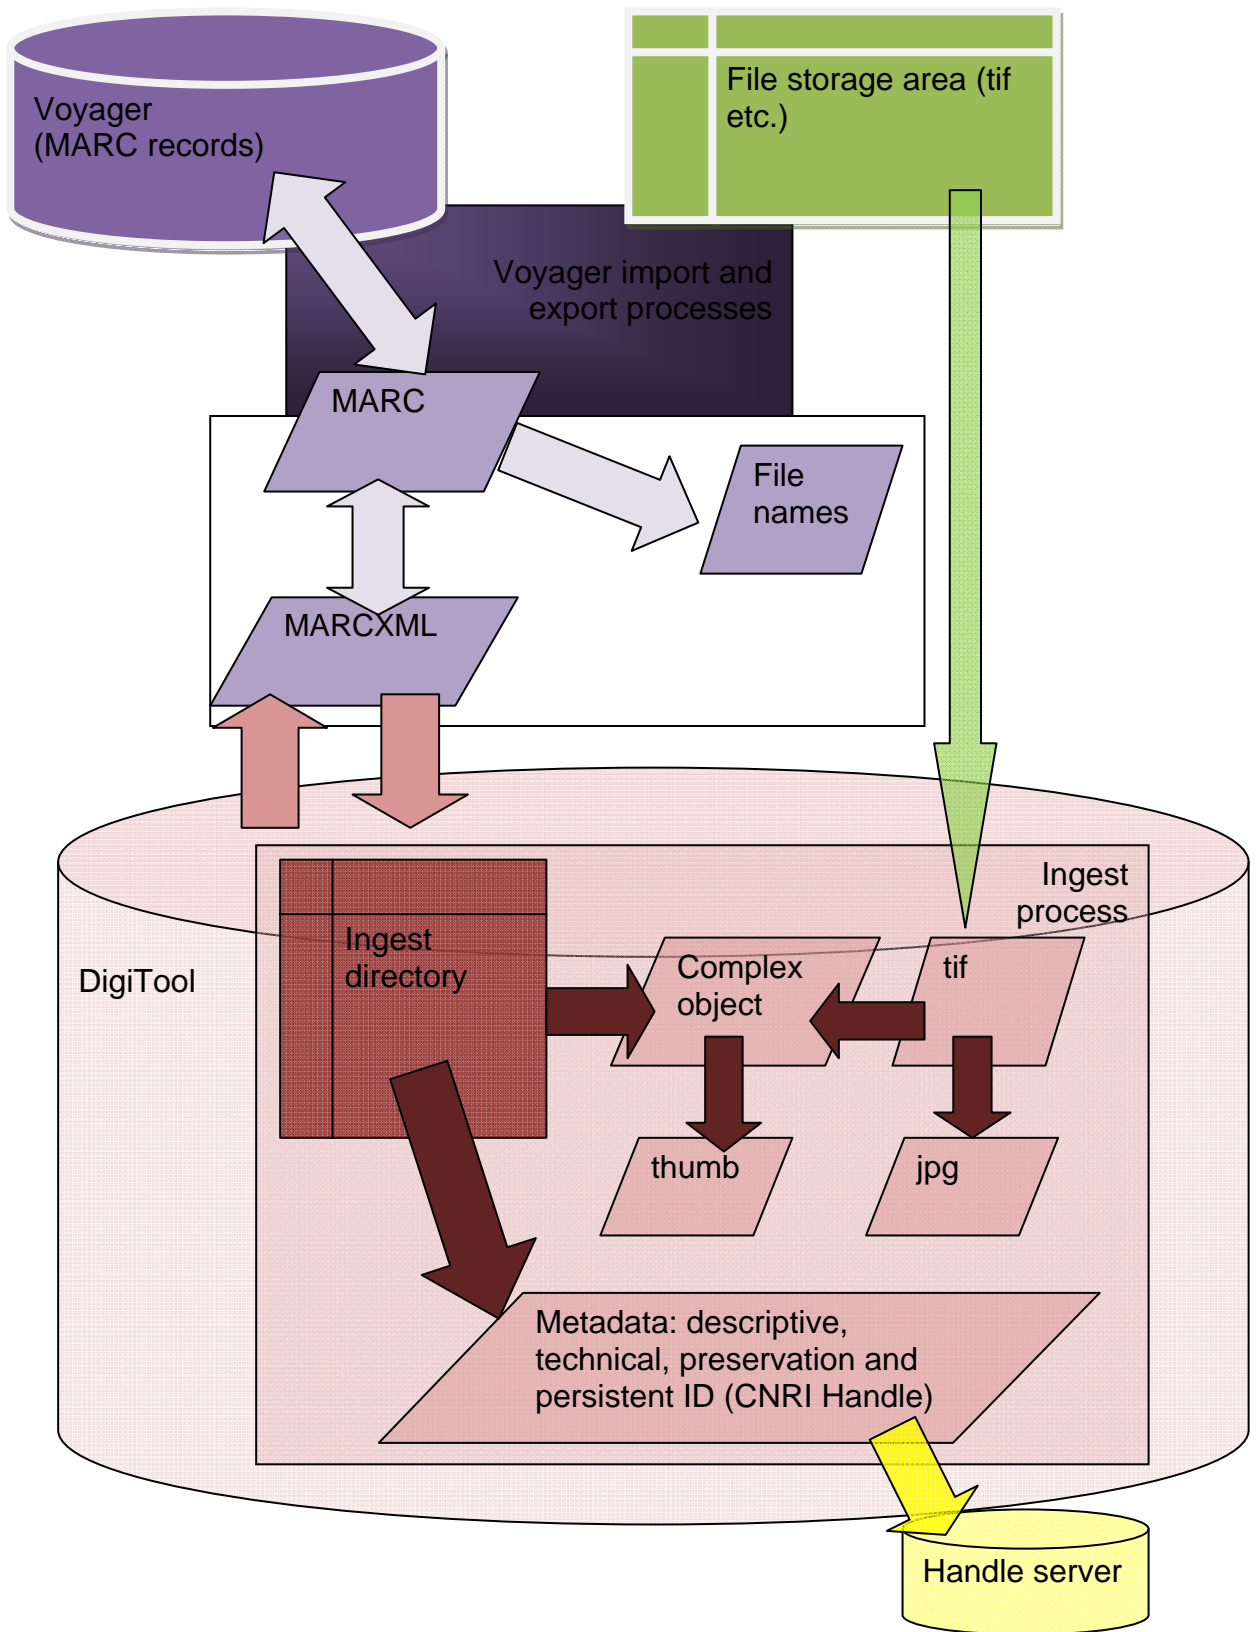
Part one: bibliographic and holdings data

How do the systems mentioned in the introduction currently use this type of data at the State Library of Victoria and how is this managed overall?

In order to cover this issue in proper detail we have produced some diagrams covering data modelling within single systems and between a number of systems. Some of these diagrams overlap between their coverage due to modelling based on implementation work and planning. However, they do give a good picture of where each piece of data is sourced from and how it is synchronised between systems.

Continued next page

Figure 1. DigiTool and Voyager



The diagram above shows the data source and how it is copied from Voyager to DigiTool. Descriptive MARC data is sourced from Voyager (at the top left of the diagram). The MARC data is extracted based on the primary key (bibliographic record number) and converted to MARCXML. All of the complex elements which need to be moved into DigiTool are then processed as a group, so that the MARCXML metadata, technical metadata, image names and objects are all ingested and transformed into their appropriate manifestations in DigiTool. This data is split into various components in order to provide appropriate labels for the objects in DigiTool. For example, the image accession number for each TIFF image is used as a label for the image, so that the user can identify which image they are viewing, in case they wish to order a higher quality reproduction from the library, or in order to ask for permission to publish the image. The title is used as a label for the entire complex object, so that, when viewing or printing the image, the item is identifiable via the overall title of the image or set of images.

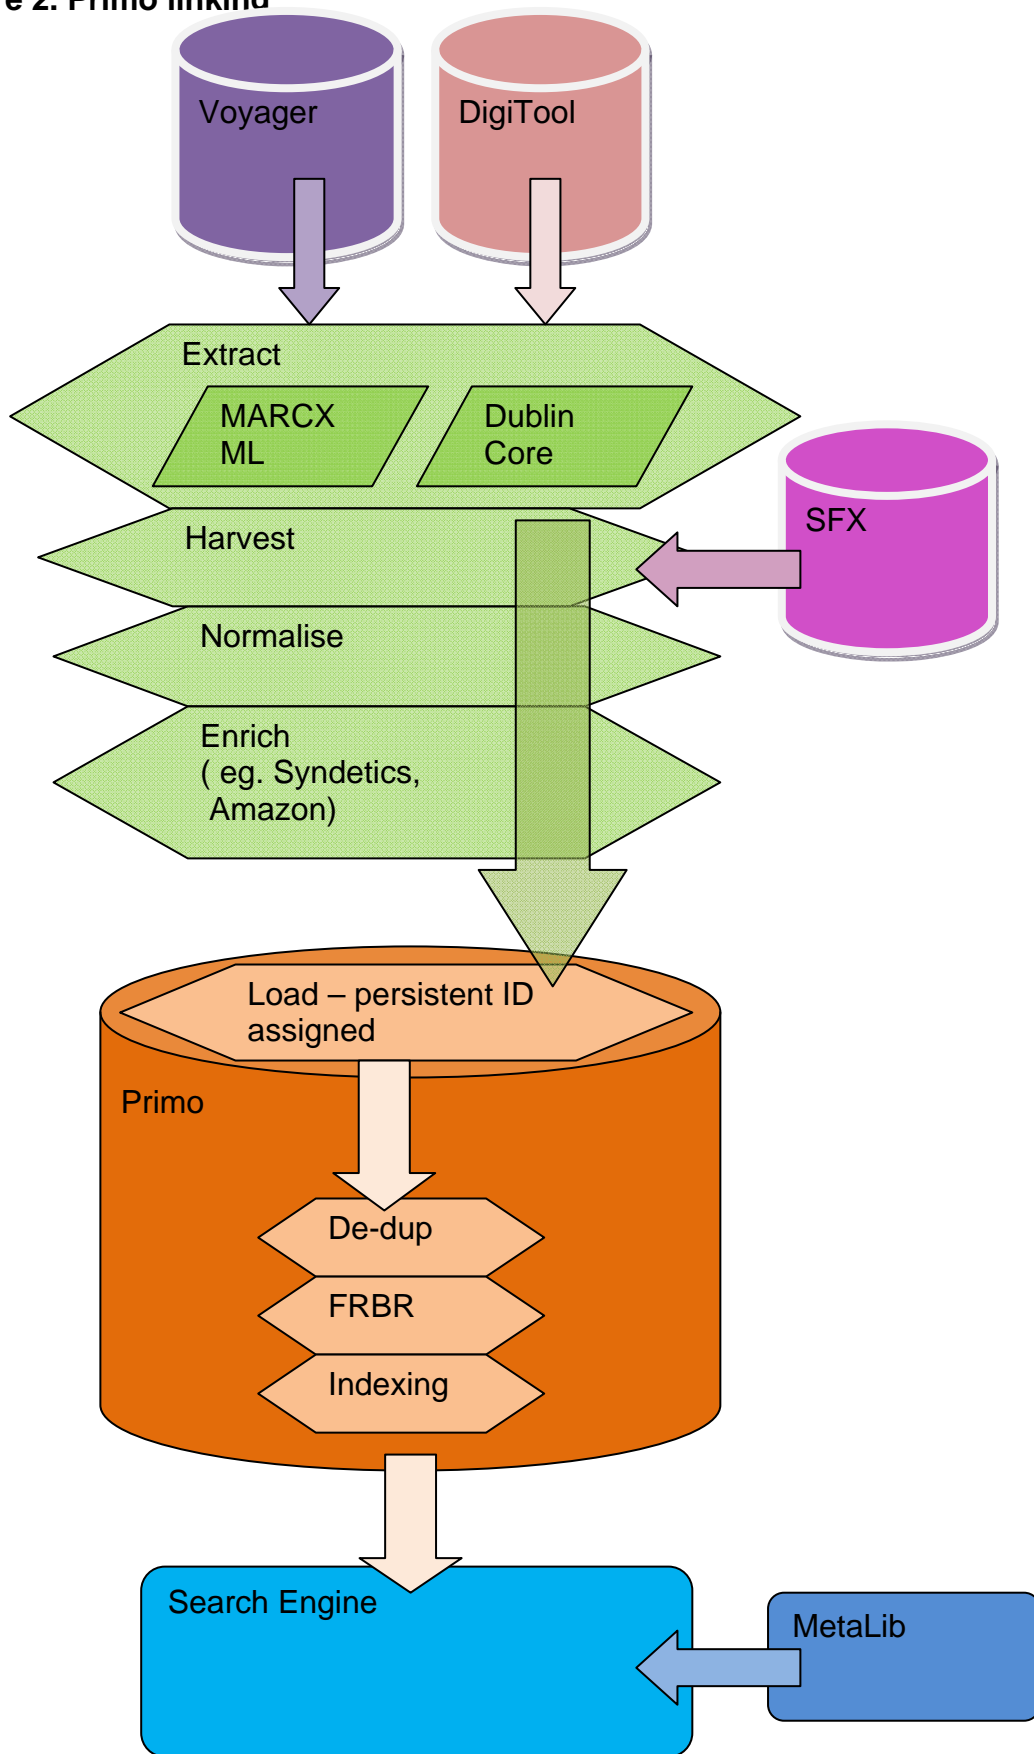
As the object is ingested into the DigiTool system, there is also a small amount of technical and administrative metadata which is assigned to the object overall and to each component. For persistent identification and external referencing of the objects, a handle, using the Corporation for National Research Initiatives (CNRI) Handle system, is assigned at the overall complex object (not at each image/component level). For visual identification of the components, a thumbnail image is also created for each part. The Uniform Resource Identifiers (URIs) for the handle and the first thumbnail image are then written back into the appropriate MARCXML field for the object, and that record is then transformed back to MARC and written back to the Voyager system once the ingest is fully completed, thereby completing the linking process between the data/resource discovery repository and the DigiTool image viewer.

If there are any subsequent changes to metadata in the Voyager system for a set of images, the MARC record is again exported, transformed to MARCXML and written back to the metadata stream in DigiTool in order to reflect those changes. These changes may be to do with corrections to images, due to errors in the scanning or due to updated information about the images since the item was scanned. There may also be corrections needed to do with the description of the object or the title and so on.

Another part of the linking between systems is the links which are made available in the DigiTool system to allow the user to move straight back to the Voyager system to view the full Voyager record (and any other appropriate holdings) or to do another search. These links also rely on the MARC metadata to be accurate and the primary key of the bibliographic record to be stable and up to date.

These transformations, extractions and updating of data between systems are not automatic. They have been carefully crafted to ensure accuracy of identification, resource discovery and viewability of items in two disparate systems, linked only by URI links. The immediate issue, which becomes obvious after setting up all of these links, is thinking about a new or better way of synchronising data that is exactly the same or very similar. The conclusion of this paper will include further discussion on this point.

Figure 2. Primo linking



The Primo system is a resource discovery and delivery interface designed to provide a Google-like experience for searchers. The system provides fast and relevant results to users by harvesting data from source systems such as library catalogues, image databases, research data repositories, abstract and index databases, full text e-journal and e-book data sources and so on.

Because this system is not a cataloguing or descriptive data input or editing tool, it is not of itself a source database for the management of descriptive data across our library systems. Therefore, the management of the source data coming into the Primo system also requires appropriate configuration to allow relevant and up-to-date data synchronisation with the source data systems. Fortunately the vendor, Ex Libris, has done a lot of the programming and design of the system to enable updating of data automatically from whatever source system you wish to point it to. However, there is still a degree of planning and management of those systems in order to maintain the integrity of the data, and particularly in this case, the display and use of the data within Primo. Because each of the data sources is quite different, and because the Primo system also manages links, thumbnail icons for this data, it is another challenge to maintain this so that the appropriate relevant data is displayed to the user at the right time.

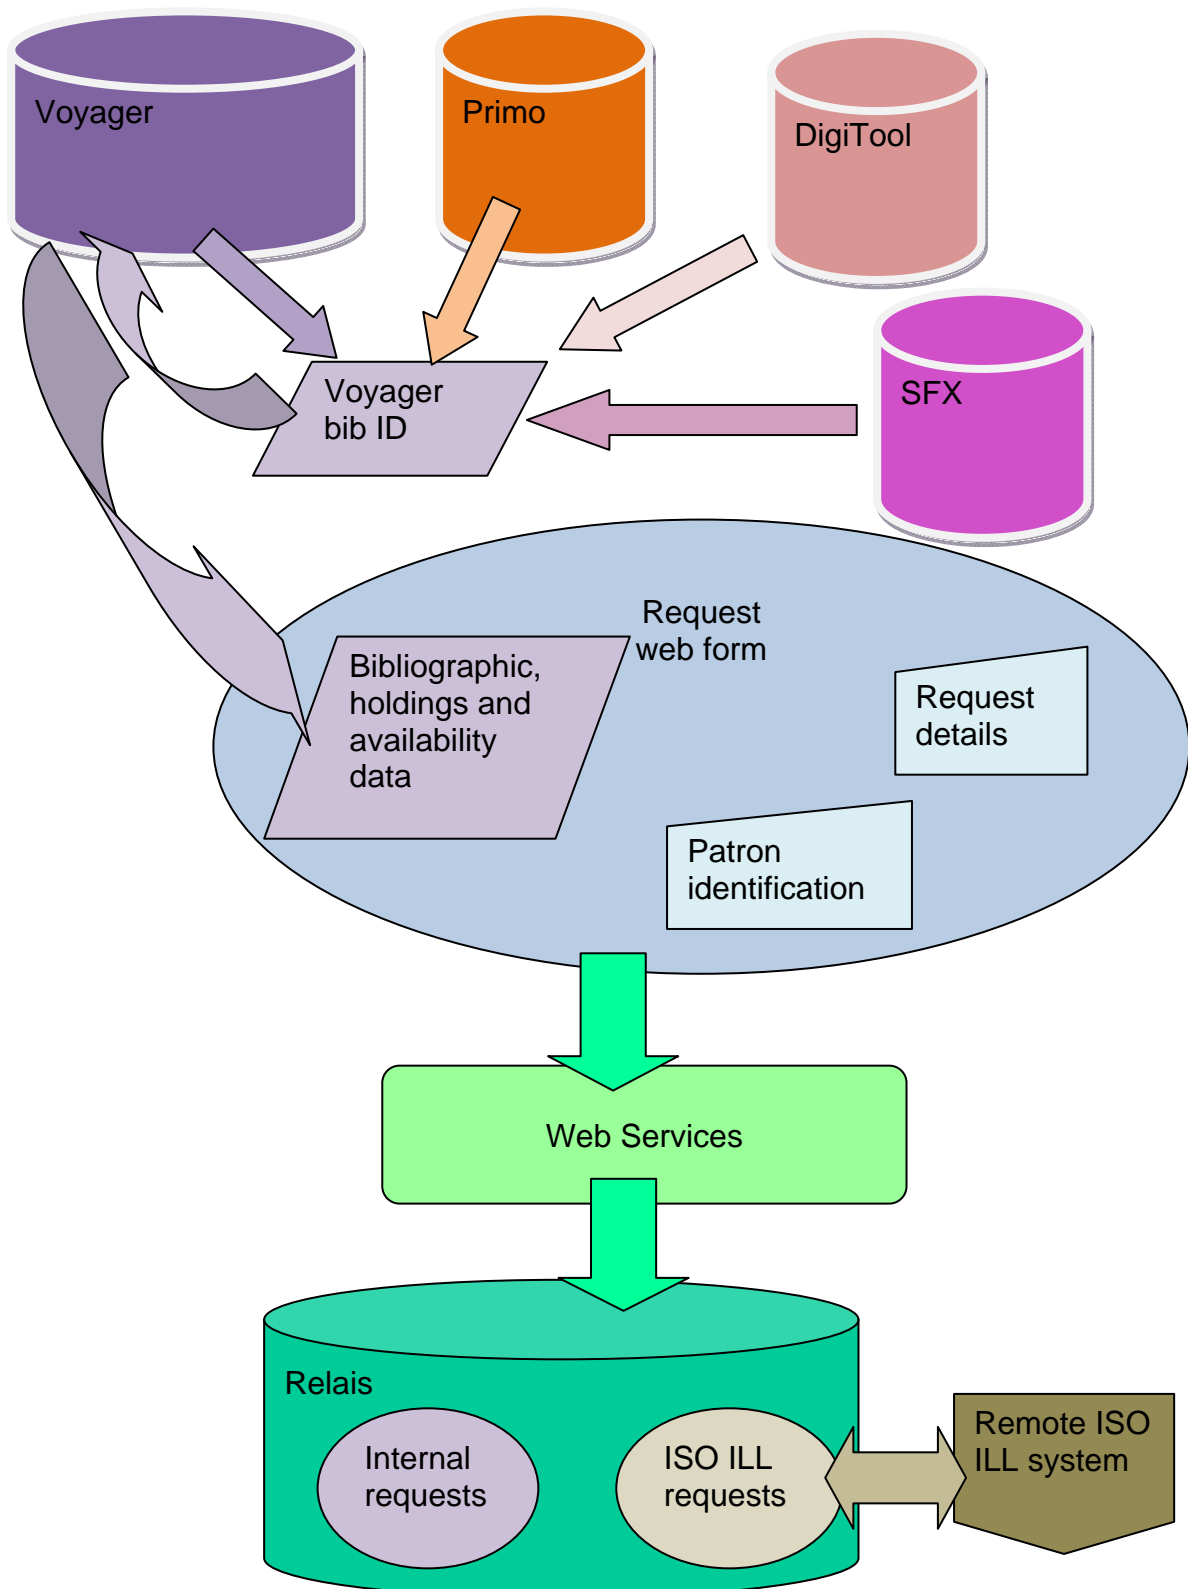
Any data in Primo needs to be managed properly at the data source. One of the issues the State Library of Victoria has dealt with is some data clean-up of the catalogue source data, as well as dealing with how the system displays and ranks online content. Primo gives preference to online versions of items, in order to provide the quickest and easiest version for the user to use, but this can often be misleading when links are not exactly what the user might want or that the State Library of Victoria might wish to show to the user. In addition, with data sources such as manuscripts and pictures, the Primo system has trouble delivering this data in an easy and relevant way to the user through the Primo interface. So some massaging of data and the interface itself within the DigiTool system has been necessary.

The Primo system uses a set of processes, called pipes, which contain new and updated data from each source system. For Voyager this bibliographic, holdings and item data in a MARCXML format (for all data types including manuscripts, images, finding aids, books, journals, e-books and so on). For DigiTool this is a set of finding aids encoded as EAD XML documents. For other data it might be OAI XML data harvested from an external source. As already mentioned, because Primo is designed to preference online versions of items (often e-books or e-journals), this is problematic for images which have been catalogued via single records for both online and physical items.

For example, there are links to and from DigiTool objects within Primo. The out of the box configuration in Primo displays external web sites, image displays and other sites within a frame underneath some of the bibliographic data. Due to the DigiTool interface being designed as a fully blown image display interface with lots of its own controls to download, zoom, etc, this meant that users really couldn't see images in DigiTool if they came from the Primo system. This entailed having to open the DigiTool links in a new window. This, however, in itself caused issues with the DigiTool interface then having appropriate links back to the search (Primo) system. So work was done to include descriptive metadata in DigiTool and links back to the detailed record and to a new search in Primo.

In terms of links to these objects, the wording for links to DigiTool needed to be massaged within a Primo 'normalization rule' in order for the links to images to have consistent wording and behaviour when the user clicks on the link.

Figure 3. Relais requesting

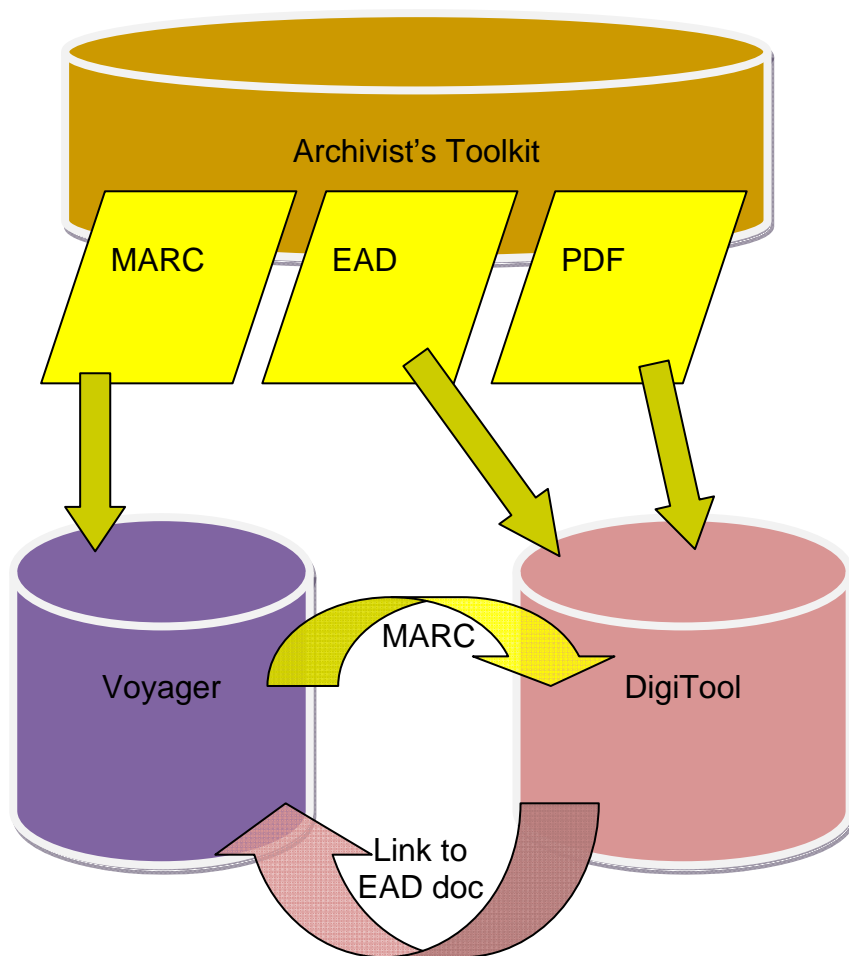


The Relais system, in terms of integration, has probably produced the biggest challenge for the State Library of Victoria in terms of the amount of planning, thinking and work needed to integrate systems which again have the same or similar metadata in them, to ensure that the Relais system receives the correct information for a user to request items or copies of items. Relais conceptually is the common requesting/linking system among all of the major enterprise library systems at the State Library of Victoria. However, it does store the least amount of data internally about items held at the State Library of Victoria. This system therefore relies totally on the correct data for the item being requested to be sent through to Relais from the other systems.

Relais has the architecture of a well written, designed and fully functioning Interlibrary Loans (ILL) system. The system supports the usual International Organization for Standardization (ISO) ILL standards for messaging between ILL systems. This method does allow the transfer of bibliographic data between ILL systems for the purpose of requesting items from other libraries, or for fulfilling requests for items from other libraries. Usually, this mainly includes standard call number, location and dates/issue/page data being requested, but there are often gaps between this data and the actual item/copy holdings data that the State Library of Victoria itself holds (for incoming requests). Because of the fact that the State Library of Victoria is not a lending library, and the number of requests annually is less than 5000, this is not really an enormous problem. The challenge for the State Library of Victoria has been to enable sometimes complex and inconsistent item/copy/holdings data across diverse collections (with diverse cataloguing data) to be sent correctly to the Relais requesting system, to enable requests to be accurately fulfilled.

As outlined earlier, because the Voyager data is the common source of data amongst all of these systems, a decision was made to use that data source for all requesting functionality from any of our systems, no matter what they were. So, for example, if a user clicked on the "Request a Copy" (or appropriately named) button in Primo, the data which was transferred into the online Relais request form was verified and sourced (via the primary key - the bibliographic record id number) from Voyager. Each of the systems, Primo and DigiTool (and possibly SFX via the SFX services menu call), all contained Voyager-sourced metadata internally which included the same bibliographic record id number from Voyager. Therefore, it was easiest for us to use the same method to go back to the same source system and pull out the data in the same way, no matter what source the request came from. This meant that only one script was needed to support the same piece of functionality (requesting of bib/holdings data into Relais). It also meant that data was delivered in a real-time manner, since the Voyager system was the central controlled data base and authoritative source where the library maintained its bibliographic descriptions of items and item/holding/copy data.

Figure 4. Archivists Toolkit



One of the anomalies in this scenario is the Archivists Toolkit database. This system was trialled in early 2009 and rolled out into production in August 2009. The system is an open-source Mellon Foundation-funded system developed in the US. This system supports the accessioning, provenance, description, management and many other functions required by our heritage areas in order to manage their large collections.

In the past, the State Library of Victoria has used a variety of mechanisms to perform these functions. The HDMS system (Heritage Data Management System) developed by the University of Melbourne's Australian Science and Technology Heritage Centre (now the eScholarship Research Centre) was used for cataloguing, accessioning and descriptive functions, and the Voyager system for basic descriptive functions, hand typed Descriptive Lists (many now in PDF format on the web) and the Oxygen XML editor to hand code Encoded Archival Description (EAD) finding aids for use via the DigiTool system.

The Archivists Toolkit, like the Descriptive Lists or EAD finding aids, represents a change in methodology where the description of collections is done outside of the main cataloguing system (Voyager). However, the Archivists Toolkit will enable export of a range of data into MARC, PDF and EAD formats to support the discovery and delivery of this data via these other systems. In this case the data flow is quite different and changes the origin of the data from focusing on the Voyager cataloguing system to Archivists Toolkit. In this case, Voyager and DigiTool are the target systems. However when the EAD document is ingested into DigiTool, the descriptive MARC metadata is also copied to DigiTool and associated with the EAD document, with links being created back to the source MARC record in Voyager or Primo.

Integration issues

One of the advantages of having homogenous data across the different systems has been to allow the primary key (in the Voyager database this is the bibliographic record number) to be used for linking to and from each system. This has worked well in many circumstances. However, one area which we were not able to control very easily was allowing the user to seamlessly link back to their original search pages by referencing the referring URL or data indicating where they came from. For example, if a user landed straight into DigiTool from Picture Australia, we were not able to configure the DigiTool system to link back to the record in Picture Australia, because we were not capturing a dynamic link within the system that would indicate the referring URL. This data is recorded within the Apache logs, but in the image/object viewer in DigiTool it is almost impossible to use this data on the fly and reference it during a user session. This is due to the viewer being written in Java libraries provided by the vendor, which are very difficult to dig deep into, modify, and control. So we had to make decisions about what we would link back to, in order for users to go somewhere to a new search. In this case, before we had gone live with the Primo product, we used the bibliographic record number for each object to link back to the full bibliographic data in the Voyager system, as well as providing a "new Search" link back to the Voyager search interface. When Primo was released to the public, we changed this link to preference a new search and a link to the full record in Primo.

As our systems become more outward reaching and our data is harvested by more search engines, research repositories and data archives, we would hope that this problem would become more recognised by vendors as a specific need and we will easily be able to link users back to their source search interface on the fly, as would be expected with most modern user interfaces.

Part two: Patron data

The issue of patron data has been much better managed within enterprise-wide systems in large institutions for a long period of time. The use of central data registries as one single authoritative source, and pointing of systems towards that data, has been driven by the need for central authentication and single sign-on, as well as security and privacy concerns. For all of these reasons, the centralised authoritative model of one single data source is really the standard in the industry. If institutions are not doing this, it is either because they are small and are able to manage separate authentication and registration systems easily, or they do not have large resources to purchase (or build) central well-managed patron registry systems. As new systems are purchased and come on board, it is essential for systems to conform with such an architecture, particularly when dealing with sensitive personal information. Almost all modern well-designed systems do conform to this model, or are able to be modified easily to fit in with the central patron data model because standards and models are well known and supported across most systems.

At the State Library of Victoria, the issue of patron data management has been a challenge for a very long time. This is because:

- patrons come and go
- our potential user base is very large
- we don't know our users until they walk in the door or start to search/request/use our systems
- our data is very wide ranging and visible outside of our institution, so a wide catchment for users is created

Our patrons also use our library in many different ways, including exhibitions, business meetings, foundation members.

The State Library of Victoria does have a policy for user engagement, which encompasses the following principles:

- low barriers to engagement (only ask for personal information where essential)
- simplicity
- persistence (patrons only need to log in once per session, regardless of services accessed)
- compliance with privacy requirements


The web site redevelopment project, SLV2.0, has an element covering authentication, which will resolve many issues outlined above, by installing a central well-managed user registration, authorisation and patron data storage system. When this system is ready and working (which we believe will be some time in 2010, all of our vendor products, as well as all of the key parts of the web site, will be configured to point to this central patron database.

In the meantime, however, we do have a patron management problem. The SLV2.0 work has identified 14 or so separate patron databases that the different areas of the

library use for separate purposes, such as marketing, communications, exhibitions, requesting and borrowing, e-journal authorisation and so on. This has been a challenge for the library to manage and is clearly an area for improvement which will be resolved with the new central patron database.

As far as enterprise-wide vendor systems are concerned, because they mostly are from one vendor, they are able to use the same patron data source, which is the Voyager patron database. This data source is used as a source of authentication, identification, as well as authorisation. The Relais system can also be programmed (with a bit of work) to point to a separate patron database, aside from its own internal one.

Figure 5. Patron data management



This patron model of always pointing to the one central data store, and using that for a variety of functions, is a very good one and lessons can be learned from this model for the design of other data models.

New Models for data management

All of the enterprise-wide systems being implemented are commercial-produced software; they have not been developed in-house. This has meant that opportunities for full-scale data modelling has been limited; given that we were dealing with established data structures, our main scope was restricted to two areas: the mapping of data elements between systems, and the creation of small scale data structures as part of in-house-developed functions to manage data. Our implementation of Relais is a good example of both of these areas.

Data mapping

We decided quite early on in the project that we would need to develop our own web-based request forms for patron requests. Our reasons were fourfold:

- We will be using the Relais software not only for traditional ILL functions, but also for provision of electronic copy of our own collection materials to our own patrons.
- Our collections are diverse and access to them is supported by complex business rules, as is delivery of material. The request form had to take into account many of these rules, for example to ensure that patrons would be directed to digital copy where it was already freely available.
- We wanted to take advantage of existing bibliographic and patron data wherever possible, to avoid unnecessary duplication of data and to minimise the amount of direct input required by the patron filling out the form.
- We offer specialised copying services and needed to be able to allow the patron filling in the form to express their requirements.

The decision to create our own form meant that we had to do some fairly complex mapping of bibliographic and patron data from Voyager Web Services and patron-specified request data from the web form, to Relais Web Services. We also had to gain a clear picture of where the data would reside in the Relais data structures, and how it would appear to the staff member processing the request in the Relais client. Some of this is still work in progress.

Data structures

In the early stages of the web request form development, it became apparent that effective session management required that we create a small database to hold patron, session and other elements from the web form in order to manage the session state for the user.

Data management issues

The problem with all of this data copying and synchronisation is an obvious one. Why has it been done this way is probably founded in the very successful Google model of pulling in data and indexing it locally to enhance search speed and accuracy. However, the solution is also an obvious one. Why don't these systems, which all use Oracle databases at their back-end, use Oracle single business integration tools to share data in a more efficient way, rather than relying on scripts to do the work? The question to be asked is: is it more efficient to transfer data on the fly from a source database using direct database tools and via indexes within Oracle, at the point of need, or is it more efficient to transfer large pieces of data in advance and index it locally, and only transfer the smaller pieces of data as required on the fly? What seems obvious, with systems that are locally housed anyway, is to use indexes and transfer data directly, without having to build new tables and indexes in another (side by side/local) system which already has tables and index structures to get at the same data.

The major resource-discovery vendors are already beginning to use a model similar to this to help them get to full-text data sources from publishers in a more efficient way. Vendors are changing from a federated search model to directly searching remote indexes, to enable fast resource discovery from remote publisher databases. This is a much better model for search efficiency, particularly when libraries and vendors have been struggling, and users have shown a lot of dissatisfaction with, the slow, clunky (so last-century) federated search model. Other significant factors have been the limitations of the Z39.50 standard in dealing with complex queries across source SGML, XML and other metadata schemas and the limits of MARC data and how that is indexed when being presented back to the federated search engine.

The announcements last year by Serials Solutions, Ex Libris and other major library vendors about securing indexed content from a very large number of vendors, in order to enable direct, fast index searching from their resource discovery systems, is a major step forward. One of the important selling points and value that universities put on their libraries, and what can be delivered to users, is scholarly published content. With vendors now being able to deliver this content much faster, this really moves the resource discovery interfaces into a Google-like simplicity for resource discovery, with a very powerful fast back-end delivery mechanism. All of the recent technology aspects of library system interfaces (which allow for faceted browsing, FRBRisation, grouping and limiting/scoping of data) is now making data much more valuable within systems other than local catalogues, image databases or research repositories.

Conclusion

The old model of copying data between servers, across systems and running things in batches is such an old way of working with data, based on an old view that systems are all separate and can't talk to each other directly.

At the State Library of Victoria, as is demonstrated in figure 6 below, just in terms of our major Library systems there are currently eleven different jobs which we are running to copy data back and forth between those systems, in order to keep the data synchronised.

Figure 6. Data maintenance jobs

2.00am	Voyager	gen_xml.sh slvdb	Extract updated Voyager records to go to DigiTool
2.30am	Voyager	postprocess_linking_tool.pl	Rewrite links to digital objects into Voyager records
Every 5 mins	Voyager	Patron registration update	Patron validation online
7.40am	Voyager	patron_online_registrations.sh	Patron online registration
6.00am	Voyager	slv_dop_phase1.sh	Extract new Voyager records to go to DigiTool
2.00pm	Voyager	slv_dop_phase5.sh	Prepare new XML files from DigiTool to write back to Voyager
10:15am, 1pm, 4pm, 8pm	Voyager	primo-exp.sh	Export changed Voyager records for Primo
7.00am	DigiTool	slv_dop_phase2.sh	Move new MARCXML records to ingest sequence in DigiTool
1.00pm	DigiTool	slv_dop_phase4.sh	Ingest process for DigiTool
8.00am	DigiTool	maintenance_meta_update.sh	Update selected records which need maintenance in DigiTool
10.00pm	DigiTool	nightly_maintenance.sh	Update records in DigiTool

We believe that library vendors (and library staff, programmers, developers, etc) should be moving away from a scripted/timed and batch-based methodology to a more seamless XML-based real-time model for the updating of data.

A newer/better model would cover the querying of indexes directly on remote systems, for all sorts of data that the library is trying to manage. As is demonstrated in figure 5 above, this would be most efficient when including local catalogue data, to simplify the data movement/modelling/scripts and other parts that have to happen currently with local data in resource-discovery interfaces.

Real time querying and direct updating of data at the time of need is the most efficient way to ensure all of our systems and services really do seamlessly connect and to integrate between them.

References

Akerman, Laura. Making it Work in 2 Places: Maneuvering between Emory University's EUCLID catalog and Primo "discoverE" system. 2009
http://presentations.ala.org/images/e/ea/Making_it_work.ppt

Boehr, Diane. On the Record: Report of the Library of Congress Working Group on the Future of Bibliographic Control, 2008.
http://presentations.ala.org/images/7/70/ALA_Future_of_bib_control.ppt

Calhoun, Karen. The Changing Nature of the Catalog and its Integration with Other Discovery Tools, 2006.
<http://www.loc.gov/catdir/calhoun-report-final.pdf>

Christenson, Heather. Integrating Information Resources: Principles, Technologies, and Approaches by Heather Christenson and Roy Tennant, 2005.
http://www.cdlib.org/inside/projects/metasearch/nsdl/nsdl_report2.pdf

Corporation for National Research Initiatives (CNRI) Handle system
<http://www.handle.net/>

Duval, Erik. Metadata principles and practicalities / by Erik Duval, Wayne Hodgins, Stuart Sutton, and Stu Weibel, in D-Lib Magazine, 8(4), Apr 2002.
<http://www.dlib.org/dlib/april02/weibel/04weibel.html>

Mak, Lucas Wing Kau. Multiple Sources, Multiple Paths: Migrating Metadata Between Systems. 2009.
http://presentations.ala.org/images/8/82/Multiple_Sources%2C_Multiple_Paths.ppt

Mann, Thomas. "On the Record" but Off the Track : A Review of the Report of the Library of Congress Working Group on the Future of Bibliographic Control, With a Further Examination of Library of Congress Cataloging Tendencies, 2008
<http://www.guild2910.org/WorkingGrpResponse2008.pdf>

Marcum, Deanna B., Response to "On the Record: Report of the Library of Congress Working Group on the Future of Bibliographic Control", 2008
<http://www.loc.gov/bibliographic-future/news/LCWGResponse-Marcum-Final-061008.pdf>

Pace, Andrew K., Dismantling Integrated Library Systems: librarians and their vendors have created a tougher world for themselves, with interoperability the only solution. *Library Journal*, 2/1/2004.
<http://www.libraryjournal.com/article/CA374953.html>

Schaffner, Jennifer. The Metadata is the Interface : Better Description for Better Discovery of Archives and Special Collections, Synthesized from User Studies, 2009.

<http://www.oclc.org/programs/publications/reports/2009-06.pdf>

Tillett, Barbara. What is FRBR? A Conceptual Model for the Bibliographic Universe, 2004.

<http://www.loc.gov/cds/FRBR.html>

University of California Bibliographic Services Task Force. Rethinking how we provide bibliographic services for the University of California. Final report, 2005

<http://libraries.universityofcalifornia.edu/sopag/BSTF/Final.pdf>