Appendix I. ReCAP Project Technology Report

ReCAP Project Technology Report

Research Collections and Preservation Consortium

By Marshall Breeding, ReCAP Technology Consultant

Revised – November 2012

As part of a project to expand the vision of the ReCAP facility from a shared storage facility to a shared collection with enhanced access to the patrons of each of the participating libraries, an enhanced technology infrastructure needs to be implemented. An enhanced technical environment to support the next phase of ReCAP might include capabilities such as:

- Inclusion of materials available to end users at any of the three participating institutions in the shared ReCAP collection through existing online catalog or discovery interfaces
- Display of real-time status of items in ReCAP, including availability for request, restrictions, available pick-up locations
- Improved services that can be embedded into the online catalog or discovery services of the participating institutions to capture and validate requests made by library users or by library personnel for ReCAP materials
- Tools for tracking ReCAP materials requested by patrons from the time that they leave the ReCAP until they are returned for refilling
- Tools to support the management of the ReCAP collection, such as collaborative collection development, automated processing of duplicates, or designation of preservation retention.

This report provides an environmental scan of the relevant technology components in place in the participating institutions and posits functionality to support the operation of ReCAP as a shared collection with unmediated patron request capabilities. The report suggests some general approaches and specific products or projects that might be considered as candidates as part of an expanded technical infrastructure.

Part I. Strategic Infrastructure used by ReCAP Institutions

ReCAP Facility

The key technology issue for the ReCAP facility itself centers on the Library Archival System developed by Generation Fifth Applications that provides a comprehensive inventory management system and that automates all of workflows carried out in accessioning new materials, handing retrieval requests, reporting statistics, and a variety of other tasks involved in the operation of the facility.

Library Archival System

The Library Archival System was developed by Generation Fifth Applications as an inventory management system based on the Progress fourth-generation language. This software tends to be called "GFA" by those using it.

Information regarding the GFA software was obtained through interviews with Eileen Henthorne (Executive Director ReCAP), Steve Gilbert (GFA President and owner) and Christine Brennan (GFA Operations Manager), as well as through other ReCAP participants.

Generation Fifth Applications (http://www.gfatech.com/) has been in business for about 30 years and develops applications for other industries besides that of library storage facilities.

Gilbert reports that the company considers itself very progressive in its approach to technology and is interested in new developments that would meet the needs of its library customers. He states that access to the database management system can be accomplished in any way that the libraries need. In addition to the batch mode that ReCAP uses for most of its work, direct mode is also available that provides more dynamic access to the data.

GFA Library Archival System is currently based on Version 10 of Progress (www.progress.com), with development underway for version 11. Most customer sites have migrated from Version 9 to 10.

Gilbert considers GFA as a research and development company. He reports that the company works in other areas, such as in telecommunications. The company owns multiple data centers. GFA was an early proponent of cloud computing concepts.

GFA follows a shared source business model. While customers do not receive the source code to the system, new functionality sponsored by one customer is shared with any other sites interested in the same functionality.

The development approach is to avoid being driven by the technology but by the operational needs of the organization.

Work is underway to expand the sub-catalog, developing a catalog in its own right. Such a catalog would be limited to inventory data and not bibliographic descriptions.

Gilbert indicates that the company would be generally open to creating additional functionality, such as service layer of APIs for ReCAP.

Functionality

The Library Archival System is designed to be an industrial inventory management system with operational requirements that differ considerably from an integrated library system. Its primary function is to automate the operations of the facility and to manage its inventory in such a way that no materials are lost.

The GFA software uses a text-based terminal interface accessed through a secure shell client (see www.ssh.com). The use of ssh uses an encrypted communications to provide secure access to the

software without exposing usernames, passwords, or operational data. All data entry is accomplished through menus and keyboard commands.

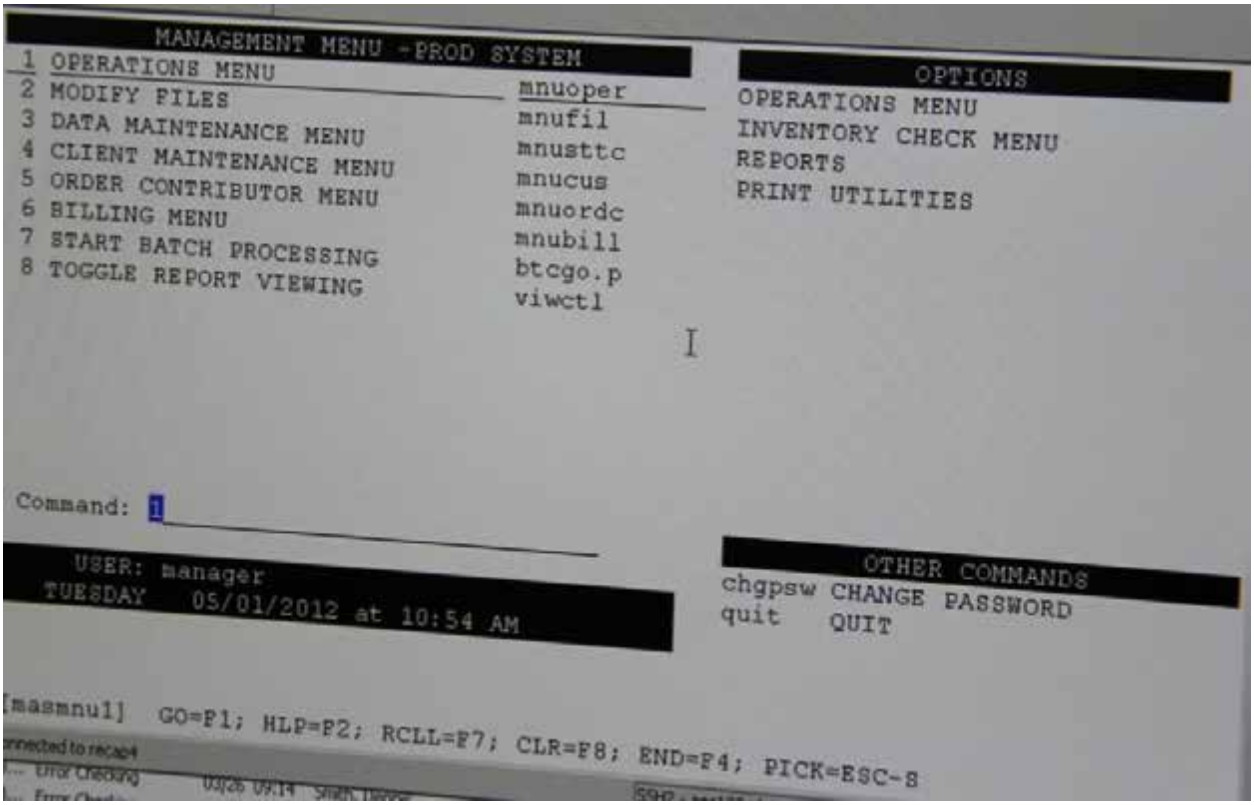

The Library Archival System and the operational procedures in place in ReCAP are designed to result in a very high level of control so that no materials are lost within the facility. ReCAP personnel have a very high degree of confidence in the system to manage their inventory and to automate their operations.

The boundaries of the control of ReCAP and the Library Archival System are rigidly defined. It controls items once they are received into the facility and accessioned into the GFA software. It exerts no control of items once they leave the facility. Libraries are responsible for tracking requested items from the point where they leave ReCAP until they are returned for re-filing. The handling of materials outside of the control of GFA leaves considerable room for uncertainty of status and availability and error.

Data Structures

Some of the data elements include a unique identifier, customer codes, location, current status information, and its transaction history.

Customer Codes are assigned to each accessioned item. These codes reflect the original institution and the collection or sub-collection. The customer code controls the business rules that apply to the item. Some customer codes represent materials that have special restrictions and may not be able to be requested by other institutions or even by some categories of users at the home institutions. Of special interest are the customer codes that can be requested by all the ReCAP participants that would comprise the shared collection.

The GFA Library Archival System maintains a detailed history on each accessioned item. This transaction log can be used for problem solving and for reports involving the frequency of use.

Within the ReCAP facility and within LAS, there is no representation of bibliographic data for any item. Bibliographic data remains in the respective integrated library systems.

System Architecture

The software has been created with the OpenEdge Advanced Business Language, commonly called Progress, a fourth-generation language that includes integrated relational database functions.

The Library Archival System, as are all products created by GFA are based on a framework called ToolPro, positioned as a fifth-generation language. The name of the company derives from its orientation toward this fifth generation language for software development.

GFA describes the LAS application as comprised of at least a hundred smaller programs, which have been assembled to meet the requirements of ReCAP. Each implementation of LAS will be different, depending on the specific programs needed for their operations and workflows.

ReCAP, for example uses a program that performs "transit bin tracking" which is used only by about a third of organizations using LAS.

LAS can be delivered to support different workflow models. The software has been implemented by the University of Oxford in the United Kingdom, for example, and their version has been configured to handle a different workflow for the accessioning of materials. Oxford uses a linear accessioning process to accomplish an accession rate of 40,000 items per day to meet their goal of processing 7 million items per year. This linear workflow involves having materials pre-sized so that a group of three personnel work together on a workbench to do accessioning and verification. In contract, ReCAP uses a parallel workflow involving three sets of operators working separately.

The GFA software is not perceived to be very open by the libraries involved in ReCAP. While there are many reports available and the ability to perform real-time queries, there are no APIs currently available to customers to extend functionality or for interoperability scenarios.

Deployment model

The GFA Library Archival System is operated on a virtual server housed at Princeton University.

No problems have been reported with this arrangement. The implementation has proven to be extremely reliable with very little down time.

The implementation can be considered a server-based system operated through character terminal interfaces accessed through secure shell clients.

Enhancement Process

While incremental enhancements are done as needed, no plans are currently underway for a major revision of the software or to consider other alternative systems.

The last major revision to the software was performed in around 2001. Since that time incremental revisions have continued to be implemented upon request from the broader GFA customer base.

Brennan indicates that there has been some interest in providing a Web-based or graphical interface for LAS. This may not be needed for operational staff of the facilities that require very efficient keyboard workflows. Use of a mouse would likely reduce efficiency. Other types of use, such as those by authorized staff from the participating libraries, or those that need to create reports would benefit from a friendlier interface.

Other future developments might include the ability to operate the software from iPads or other tablets. This capability would provide support for lower-cost devices for real-time access from within the storage modules, when wireless network access is already available.

Eileen Henthorne indicates that some of the desired enhancements that would benefit ReCAP operations might include the ability to do SQL-like queries and to have more capabilities for ad-hoc reports. But from the perspective of operational efficiency, the software in its current state is believed to work quite well.

Customer base

GFA Library Archival System finds use in most of the major high-density, non-robotic library storage facilities. Major installations include the Harvard University Depository, Cornell University Library Annex, and many others are listed on the GFA Web site. The system has recently been implemented to support Oxford University Bodleian Library's Book Storage Facility in Swindon, representing the largest installation to date. The Clancy-Cullen facility in upstate New York, with which NYPL contracts for the short-term transfer and storage of their materials, also uses the GFA system.

Large-scale library repositories represent a very small niche for automation support. While these facilities may have some characteristics in common with other types of warehouse operations, they have very specialized needs. There are not currently other products that directly compete with LCA. The only competition is a conceptual one where some libraries have chosen to use their integrated library system to manage the inventory of their storage facility.

There is not a formal user group organized for the institutions that have implemented the GFA Library Archival System, but there have been some informal meetings.

Cost and Licensing

The Library Archival System is licensed according to the number of staff clients, or "seats" that can access the system simultaneously. The initial cost of the software was around $100,000 [need specific figures] and the annual maintenance is approximately $16,000. The maintenance includes problem support as well as access to new versions and enhancements.

Observations about Library Archival System

The number of large-scale repositories using the GFA software reflects its dominant position. It does not seem likely that other applications will emerge from competitors given the number of potential clients would be limited relative to research and development costs. The GFA software has been developed

primarily for the operational needs of the repositories. Other concerns such as interoperability with discovery services and workflow management systems have not previously been part of the development strategy. The extension of the software for these purposes would not only benefit ReCAP, but would likely be of interest to other institutions.

ReCAP Functional Workflows

All tracking in the ReCAP facility is accomplished through barcodes. All materials submitted to ReCAP must have barcodes correctly positioned. The rapid processing off material depends on the consistent placement of barcodes. All containers involved in the process have unique barcode numbers such as bin, transportation totes. Each shelf in the storage modules has a barcode.

Figure 1Barcodes identify each shelf in the ReCAP storage modules

No alternative identification technologies, such as RFID, are used by ReCAP.

One of the key principles of high-density storage facilities involves making efficient use of storage space by organizing materials according to size. ReCAP stores items in trays according to factors such as size and the owing institutions.

No consideration is given to bibliographic details or classification schemes. One of the key issues for ReCAP and the GFA software is that associations between inventory data and bibliographic data must be handled through a hybrid system involving the ILS of each library.

The accessioning process begins with sizing items and assigning them to size-related categories. A given tray is filled with materials of its designated size classification. Each tray is barcoded with a unique identifier and each item placed in it is scanned and registered accordingly to that tray.

Figure 2Sizing materials as part of accessioning process

The GFA system assigns a location for the bin based on its size and available space of that category. Configuration parameters and a data representing all available space in the facility drive the system to assign shelf locations.

All trays are verified by re-scanning to ensure that all items are present.

Figure 3 Materials of the same size category go into each tray

The accessioning process is designed to ensure that any given item can be definitively located, including the module, aisle, shelf location, and bin within the shelf.

The totes used to ship materials to the libraries are also barcoded and tracked with an optional module to help ReCAP know which totes are at a given library.

Transmission of requests

Requests from the libraries are currently managed through the exchange of data files in scheduled batch operations. Each library will create a file of requested items. The file is created through scripts at each institution that assemble the accumulated requests received during that period.

Files are structured in comma delimited format, with quotes surrounding each data element. The following data elements are required in files submitted to ReCAP:

Field	Required?	Content*
Request code	y	REQI, ADDI, DELE
Item barcode	y	
Priority code	n	I
Patron name	n	
Item title	n	
Pick-up location	y	

Default pick-up location	y	
Item call number	n	
Item author	n	
Item vol/part/date	n	
Patron email	n	
Other Patron info	n	
Article vol/part/date	n	
Article author	n	
Article title	n	
Start page	n	
End page	n	
Other identifying info	n	
Notes	n	
Type of delivery	y	PHY, WEB, ARI

Table of data elements provided by Princeton University Library

ReCAP sends back to each library a file that reflects errors and the status of requests.

Each of the participating institutions performs updates to the status of items within its ILS according to these transactions. A successful request to ReCAP, for example, will cause the owing institution's ILS to flag the item as unavailable. Status updates are sent to the owning institution for any request of its items, including those made by other institutions or through interlibrary loan.

Rather than gathering requests into a single daily batch file, other procedures have been tried in the past that included processing individual requests as they happen. This method proved to be problematic.

When libraries receive materials from ReCAP that have been requested by patrons, they use their local ILS to manage notification and check-out. A typical workflow would involve a staff member using the holds management feature of the ILS to associate the item with the patron record and to issue a notice that it is available. The item would be checked out to the patron through the circulation module.

This workflow may need to be adjusted when patrons have access to ReCAP materials from the other two libraries that do not already exist within their ILS. Some process of automatically moving bibliographic and item records for requested items into the local ILS for circulation would facilitate the fulfillment of these materials to patrons.

Another set of workflows applies to the return of requested materials from patrons, through their libraries and ultimately to ReCAP for refilling. The priority of refilling materials back into the storage modules is naturally not as high as satisfying requests. A significant number of materials may be backlogged awaiting refilling. Materials in this state receive appropriate status updates and modifications of the GFA software has been made to support retrieval, including the ability for ReCAP personnel to know the specific location of items among the refilling backlog.

Figure 4 ReCAP materials returned awaiting refiling

Interlibrary Loan

In addition to the requests directly made from each library, ReCAP supports interlibrary loan requests. If an interlibrary loan request is made for an item held in ReCAP, an expedited process has been developed to allow it to be processed directly from the facility rather than having to be first retrieved by the library. ReCAP maintains an interlibrary-loan workstation for each of the institutions and its personnel fulfill these requests daily.

Electronic Document Delivery

A substantial number of requests can be satisfied through electronic document delivery. For book chapters or journal articles, ReCAP personnel pull the items, scan the article or book chapter, with direct delivery to the requestor.

To satisfy these requests, materials are retrieved from the storage modules and delivered to an operator that scans the requested pages. The electronic file is then submitted for delivery through the document delivery management application, using the Ariel document delivery software provided by Infotrieve.

Figure 5Scanning documents for electronic document delivery requests

Duplication issues

One operational issue for ReCAP involves duplication of materials. The ReCAP facility currently includes some duplication of items across the three institutions. There is some desire to reduce the amount of duplication, but this is more of a library issue, in that they pay according to the number of items accessioned than an operational issue for ReCAP.

Collection analysis is underway to help identify duplication of materials among the participating institutions. It may be helpful in the future to have enhanced discovery tools for library personnel that span the entire collections of the participating institutions to assist in developing collections without accidental duplication.

ReCAP Governance

The ReCAP facility is governed through a board comprised of the University Librarian and Chief Financial Officer from each of the three institutions that comprise the consortium. ReCAP is separately incorporated as a non-profit. The personnel of ReCAP are employees of Princeton University.

Eileen Henthorne serves as the ReCAP executive director and reports to the ReCAP Board of Directors.

The scans made for electronic document delivery are not retained due to copyright concerns. [It might be worth investigating whether these scans could be transferred to HathiTrust. It seems that an automated process could be developed to perform OCR, associate metadata, and other tasks to turn them into reusable digital objects that could be used under the same constraints as other items in HathiTrust. The cumulative amount of material collected is substantial it would be a positive contribution to find ways to gain additional value from them.]

Technology Issues

One of the key problems in the new environment will involve maintaining the status of any given item stored at ReCAP. Each of the individual integrated library systems indicates when an item has been transferred to ReCAP and its availability status. When the library requests its own materials, it naturally has the data needed to update that status.

When another library requests an item, ReCAP also sends a message so that the owning library can update its status in the ILS.

The lack of a modern service layer of application programming interfaces (APIs) is a key issue for this project. The new system needs to be able to interrogate the GFA software in real time in various scenarios regarding its current status and availability.

Access to the LAS through only through the menu-based character online interface and through the batch processes based on exchange of text files constrains how the software would fit into a modern service-oriented environment. While the character interface provides efficiency for internal operations of facilities such as ReCAP, the availability of a set of APIs implemented through REST services would help integrate the system into a broader technology environment for the broader consortium.

New York Public Library

New York Public Library currently has 3.5 volumes stored in the ReCAP facility, with plans to transfer a significant number of additional materials as part of its Central Library Project. This renovation project not only involves the transfer of many volumes out of the building, but also has a very high level of public attention, increasing the importance of service improvements gained through the next phase of the ReCAP project.

Many differences obtain between NYPL and the other two ReCAP institutions. Different use patterns apply to its research collection with materials primarily used only in its reading rooms while Princeton and Columbia allow most materials to be checked out to their users. These differences in policies will need to be taken into consideration as workflows are developed and as technology components are constructed in support of these workflows.

NYPL along with Columbia University and New York University formed a new cooperative relationship, called the Manhattan Research Library Initiative, or MaRLI. Through this new partnership, NYPL users otherwise not affiliated with Columbia or NYU will be able to arrange borrowing privileges at those

institutions. Reciprocal arrangements would also allow PhD candidates and faculty members of Columbia and NYU access to NYPL's research collection. MaRLI will also involve some degree of collaborative collection development among these three institutions. MaRLI has only an indirect on ReCAP

NYPL also participates in the IDS Project, a resource-sharing cooperative among a number of institutions in the state of New York. See: idsproject.org).

All three ReCAP institutions are involved as partners in the Google Book mass digitization program. NYPL announced its participation in December 2004; Columbia joined in December 2007; Princeton joined in February 2007.

Likewise NYPL, Princeton University, and Columbia University all belong to HathiTrust, a collaborative digital repository, which includes the library scans from the Google Project in addition to materials from other digitization projects among the members.

Automation Environment

The New York Public Library currently has an automation environment based on the Millennium integrated library system from Innovative Interfaces. Millennium was placed in production use in 2009 following an implementation that for the first time combined the NYPL Branch Libraries, migrating from a Dynix Classic system, and the NYPL Research Libraries, migrating from an existing Millennium system. The contract with Innovative to implement Millennium was announced in March 2008.

There is not a single library system or automation environment that spans all the boroughs of New York City. Queens Public Library uses a VIRTUA system from VTLS; Brooklyn Public Library uses a separate Millennium implementation.

NYPL makes use of multiple modules of Millennium, including Cataloging, Circulation, Acquisitions, Circa (wireless circulation and inventory), the WebPAC PRO online catalog, and Ecommerce (electronic online collection of fines and fees). Encore was originally licensed from Innovative, but it has since been replaced by BiblioCommons.

NYPL includes a large number of neighborhood branch libraries in addition to the research library. The BiblioCommons discovery interface appeals more to the type of use associated with the branch libraries. BiblioCommons, even though it is currently the default search presented on NYPL.org will not be positioned as the interface for the research libraries. The "classic catalog" based on WebPAC Pro from Innovative Interfaces will be the main search interface designated for the users of the research libraries.

NYPL does not have specific plans in place to replace Millennium, though there is some interest in moving to Sierra. The migration to Sierra could come as early as late 2012 or at probably within 2013. (At the August 21, 2012 meeting, NYPL indicated that it is planning for a transition in late 2012 or early 2013; by the September 11 call the implementation had been delayed until about June 2013.)

Depending on the timeframe for the implementation of a new technical environment for ReCAP, support will need to be in place for Millennium even if Sierra or another system will eventually be implemented for NYPL.

Millennium and Sierra

Innovative Interfaces developed its Millennium ILS beginning about 1997 as the forward path for its INNOPAC ILS. Millennium was developed with a modular client/server architecture offering staff clients created in the Java Swing framework and a Web-based online catalog. INNOPAP relied on text-based telnet clients and included very mature functionality, especially for technical services operations. Millennium carried forward that functionality into a graphical environment.

Millennium offers a very rich set of functionality across the main areas of library operations, including cataloging, acquisitions, serials management, and circulation, implemented through staff clients specialized for each module.

The initial implementation of Millennium included the use of the Encore discovery service developed by Innovative Interfaces, which has subsequently been replaced by BiblioCommons.

An important characteristic of Millennium involves the high degree of control that Innovative exerts on the system. It is a proprietary system with only limited avenues for organizations implementing the system to gain programmatic access to its underlying data and functionality.

Innovative offers some packages of application programming interfaces that libraries can license:

1. Patron API – procedures related to authentication, to query the patron database for records, to query or update patron record data, and to set or clear fines.
2. Item Status API – procedures related to the query of items, updating status, and performing check-outs
3. My Millennium API
4. NCIP – NISO Circulation Interchange Protocol used primarily in resource sharing environments, but also designed for self-service and other applications requiring access to patron and item functions.

Sierra represents Innovative's efforts to create a next-generation automation environment based on current technology while moving forward all existing functionality. The following excerpt from *Smart Libraries Newsletter* summaries some of the relevant characteristics of Sierra:

> The launch of a new system affords Innovative the opportunity to take advantage of technology components, architectures, and methods consistent with the times, and hopefully to anticipate what features will be popular and necessary in the future. Innovative has chosen the service-oriented architecture, open source database and indexing components, RESTful web services and APIs, and engagement with library developers as some of the key elements of its new technology strategy.

Following a service-oriented architecture, the new Sierra platform has been constructed in four layers. The foundation database layer will make use of PostgreSQL as the transactional database engine and Lucene for indexing to support search and retrieval operations. Use of these open source components stands in distinct contrast to Millennium, which uses proprietary database and indexing technologies developed by Innovative, or the commercial Oracle database. The database layer connects to the rest of the application through a data access object layer using a component called Hibernate which lends database independence and the ability to maintain persistent transactions through workflows executed at higher levels. Since the database layer also supports standard SQL, third party tools such as Crystal Reports can be used to create reports or other data extraction or manipulation operations. Use of the open source PostgreSQL also results in significant cost savings relative to commercial database engines.

A services layer implements the business logic that represents all of the functionality of the system, including the detailed tasks and workflows involved in ILS modules such as cataloging, circulation, acquisitions, a set of services for managing electronic resources, and another others for discovery and delivery of content. The services of this layer are exposed to higher-level applications through SOAP wrappers.

A set of new Sierra applications sits on top of the services layer. The heart of this layer will be a new Sierra App that implements the staff functionality of the system through a unified, non-modular approach. While Millennium offered specialized clients for each of its functional modules, Sierra delivers all functionality through a single application, avoiding the need to switch among modules depending on the task at hand. This application layer will also include components to deliver bundles of API's to support a variety of external interactions, all delivered through RESTful web services (Representational State Transfer). A top presentation layer will operate above the application layer, including the client to the Sierra application, web and mobile public interfaces. This presentation layer would also include third-party applications built on top of the published API's, interfaces to social media applications, or other end-user applications that might be created.

The new Sierra platform will be offered as software that can be installed locally in a library or consortium and will also be offered through software-as-a-service, hosted in a cloud infrastructure.

[Breeding, Marshall. “Innovative Interfaces to launch Sierra: a new generation automation platform.” *Smart Libraries Newsletter* Vol 31 No. 5. May 2011.]

NYPL Public Interfaces

NYPL offers patron access to its collections through both the WebPAC Web-based online catalog module of Millennium and through BiblioCommons.

Classic Catalog: http://catalog.nypl.org/

BiblioCommons: http://nypl.bibliocommons.com/

BiblioCommons

NYPL implemented BiblioCommons as its primary discovery interface in June 2011. BiblioCommons became the default catalog interface for NYPL beginning about September 2011.

BiblioCommons follows a modern service-oriented architecture and has been deployed as a multi-tenant software-as-a-service environment. The service is hosted on BiblioCommons servers and carries their domain name in the URL (nypl.bibliocommons.com).

The following excerpt summarizes the architecture and technology of BiblioCommons:

> Like the majority of existing discovery interface products, BiblioCommons operates separately from the underlying integrated library system. Like the generic discovery interface model, data are harvested from the ILS and used to populate a separate search and retrieval environment. Products like Primo, Encore and AquaBrowser each harvest the metadata from the ILS and other local collections into an instance of the software specific to the implementation of a library or consortium. BiblioCommons takes a fairly radical departure in that data from the ILS of each participating library loads into a centralized site. From the perspective of the patron, the library may scope the search to a given library or region, but the fundamental concept of BiblioCommons involves broadly shared data. In addition to harvesting basic bibliographic records, BiblioCommons harvests holdings and item-level data as well as authority records. Even though BiblioCommons relies on a shared bibliographic database, it preserves and indexes any locally created cataloging.
>
> One of the key issues with discovery interfaces is the way that they overlap and interact with the underlying ILS. BiblioCommons shares the concept of harvesting and synchronizing data describing the collection from the ILS, but into a collective service rather than library-specific implementations.
>
> Some of the discovery interfaces tap into the online catalog features of the ILS for item-specific displays and services, such as placing a hold. BiblioCommons completely replaces the online catalog of the ILS, managing all aspects of the way that the patron interacts with the collection. Its emphasis on the patron and social interactions require a much more sophisticated approach than a simple hand off to the patron services functionality built into the online catalog module of the ILS.
>
> BiblioCommons involves extensive use of patron data. It does not harvest the patron records from the ILS, but as patrons register on BiblioCommons, they are validated against their patron record in the ILS. All social features of BiblioCommons can be invoked only after the patron authenticates with the library-assigned username and pin. This requirement for authentication engenders a more trusted social environment.
>
> The interactions between the ILS and BiblioCommons take place through a Web services layer which supports synchronization of the collection data as well as real-time interaction needs for current item status and patron requests.

> Like many of its competitors, BiblioCommons relies on Apache Lucene and SOLR as its core search engine technology. Other components include PostgreSQL. Most of the server internals have been implemented in Java, including the data integration and service layer. The Web application layer was developed with Ruby on Rails, a very flexible programming environment that has recently seen widespread use in the open source community. BiblioCommons implements a service-oriented architecture, with an API made available through Web services. All communication between the BiblioCommons' own Web application and the internal server applications operate through the REST (representational state transfer) API, which is also available to the library for any custom applications it may choose to implement.
>
> Communication between the ILS and BiblioCommons takes place through a software application that resides within the library's technical infrastructure. This connector manages the transfer and synchronization of bibliographic, item-level, and authority data between the ILS and BiblioCommons. This application monitors the ILS for any changes so that data can be synchronized in as close to real time as possible. The connector also handles the interactions with the circulation module and patron data needed as users make requests through BiblioCommons that involve the local ILS. The connectors involve programming specific to each ILS product, taking advantage of any API's that might be available as well as standard library protocols. BiblioCommons has been designed to operate with any ILS, though initially the connectors have been completed for SirsiDynix Horizon and Symphony; development of connectors for Innovative's Millennium, the open source Evergreen ILS and other ILS products are underway.
>
> [Breeding, Marshall. "BiblioCommons prepares for the next phase of roll-out." *Smart Libraries Newsletter* Vol. 29 No 8. August 2009.]

In addition to the standard licensing fees associated with BiblioCommons, NYPL made a $1 million investment in the company in support of its ongoing development.

BiblioCommons offers a complete set of discovery layer services. Its implementation requires harvesting and synchronizing bibliographic and holdings information from the NYPL Millennium installation, as well as a variety of real-time functions.

The real-time interactions between BiblioCommons and Millennium take place through intercepting and parsing HTML pages as delivered through the WebPAC servers and through calls made through Patron Web Services API toolkit licensed from Innovative by NYPL.

Bibliographic data are harvested from Millennium and loaded and indexed by BiblioCommons. The search technology on BiblioCommons differs from that of Millennium, presenting differed relevancy rankings.

NYPL along with Boston Public were development partners with BiblioCommons in the creation of new functionality to fully integrate e-book discovery and lending within the discovery environment without a wholesale hand-off to Overdrive's platform. This project was launched publicly in Jan 2012.

ReCAP requests from NYPL

Until recently patron requests for materials from the ReCAP facility have been based on handwritten paper request slips mediated by NYPL personnel.

The library has recently implemented a request form allowing patrons to make ReCAP requests online. This form has currently been implemented through the Millennium WebPAC and has not been implemented in BiblioCommons, the end-user interface used by most patrons.

Figure 6ReCAP request form available through NYPL WebPAC Pro interface

The Web-based request form requires the patron to enter the barcode number from their library card and their 4-digit pin. The Web form gathers all of data needed to submit the request to the GFA system.

The NYPL web site offers an information page for patrons explaining "Access to Offsite Collections:" http://www.nypl.org/help/get-what-you-need/access-offsite-collections

> Catalog Status: "If the item has the message ADV REQUEST, the item is in remote storage. A request for this item must be made 24-48 hours in advance. The item may not be taken from the library and it is not available for placing holds."

A courier service makes deliveries from ReCAP to NYPL on Monday through Saturday. The Saturday service was added in recent months to improve service.

The ReCAP request form is currently available only through the WebPAC Pro interface. There are plans to eventually offer a similar request form through BiblioCommons. This task will require additional programming by BiblioCommons development personnel. BiblioCommons to date has not been willing to commit to the development of NYPL-specific customizations needed for the support of the ReCAP request form.

NYPL uses the ILLiad software from Atlas Systems to manage its interlibrary loan requests (www.atlas-systems.com/illiad). As with the other three ReCAP institutions, an ILLiad workstation is maintained at ReCAP and operated by its staff to satisfy ILL requests from its collections held in the facility. Requests for articles and book chapters are fulfilled through scanning and electronic document delivery.

The fulfillment process will need to recognize the different policies that apply to ReCAP materials. Some materials may not be available, for example, for loan through MaRLI. A broader range of materials may be requested by NYPL patrons for use on-site. A key part of the shared ReCAP collection will involve identifying what groups of materials, as identified by ReCAP Customer Codes, will be in the shared collection, and to develop a policy matrix of patron categories and customer codes that indicate authorized request scenarios.

NYPL staff indicated that the need for a more robust system for handling ReCAP requests. They report a fairly high number of requests that do not get fulfilled within expected time frames. Readers complain that they do not receive materials within 24 hours. Staff reported "ReCAP has been broken for some time now." No specific statistics are available on the performance of the timely fulfillment of ReCAP requests, but perceptions reflect that much improvement is needed. The problems tend not to be related to the ReCAP facility and the way it process materials, but more with incomplete or inaccurate request data or in the opportunities for error in delivery and patron notification.

The problems with timely fulfillment of ReCAP requests have been attributed primarily to inaccurate information associated with the request. ReCAP can only process requests with valid data for the item, patron, and pick-up location. Requests that include invalid pick-up locations are rejected. From ReCAP's perspective, the accuracy is 100 percent, but there are many opportunities for error outside the boundaries of ReCAP internal operations and the GFA software. From the perspective of the libraries, ReCAP requests have a lower rate of successful fulfillment than desired. The high-profile nature of the Central Library Plan increases the importance of introducing technology and processes that ensure successful fulfillment of ReCAP requests within expected timeframes and to eliminate failures.

Requests currently come in through a variety of channels, including call slips filled out by patrons in the library, e-mail requests, by telephone, and by the recently launched form integrated into WebPAC Pro. ReCAP requests are not directly managed through the holds system of Millennium. There have been ongoing problems with patrons not receiving materials within the expected time frames. Especially when the requests take place through the hand-written forms, there are many opportunities for error.

Another concern involves the expectation that high-use items will be brought back from ReCAP and returned to the on-site collection. Questions were raised about whether the concept of the shared collection supports this expectation. If an item owned by NYPL and another partner is de-duplicated within ReCAP, can NYPL still request its return?

While not directly related to ReCAP, NYPL also works with Clancy Cullen to manage the transfer of materials out of its decommissioned buildings and to stage materials out of the Stephen A. Schwarzman Building until the additional ReCAP modules are complete. A large number of materials are held in the storage facility Clancy Cullen maintains in upstate New York. Clancy Cullen also uses the GFA system to manage its inventory. When the ReCAP modules are completed, these materials will be transferred from Clancy Cullen.

Princeton University

Princeton University currently has just over 2 million items housed in ReCAP plus an additional 280,000 items from the Princeton Law Library.

With close proximity to the ReCAP facility, Princeton uses its own personnel to transport materials to and from ReCAP.

Princeton University Libraries share the vision of making ReCAP materials more easily discoverable and easily requested by the patrons associated with all the partner institutions. Princeton patrons should be able to place requests for ReCAP materials without staff mediation.

Some of the functionality that staff would expect would include tools that help with the process of de-duplication of its materials relative to the ReCAP collection. This would apply retrospectively for materials already owned that might be duplicated in ReCAP. More importantly, prospectively there is an interest in developing collections in consideration of the materials in ReCAP to avoid unintentional duplication. In addition to any patron discovery tools, there should also be a staff interface developed to support selectors in making decisions about new acquisitions, transfers to ReCAP, or withdrawals.

Other information mentioned that would be helpful to selectors would include availability of materials in HathiTrust, through the Center for Research Libraries

Both for library patrons and for staff use, there is interest in a federated index of the bibliographic information of the three institution's ReCAP holdings. A federated index of the three institution's entire collections would be useful for collection development, but may not be within the scope of the project.

Automation strategy

Princeton University currently uses the Voyager integrated library system from Ex Libris. This system has been in place since about 2000 when it converted from NOTIS.

The Voyager ILS includes a variety of APIs that enable its integration with external systems.

> According to Duong the WebVoyage Web Services design is characterized by stateless, atomic transactions and consumer/provider (client/server) architecture [1]. Each web service requires an OPAC instance to access the Oracle database. Ex Libris designed these Web Services for the following rationales: the interoperability between different systems, the integration between different applications, to enrich customer functionality, and to create a federation of resources and to facilitate faster time for software development cycle. Detailed Information regarding these services is outlined in the collaborative Web site:http://www.exlibrisgroup.org/display/VoyagerOI/XML+Over+HTTP+Web+Services. Web Services include search services, media services, patron requests services, etc. Duong (2009) further identified the process of setting up the web services. The configuration file on Voyager server is vxws.ini. OPAC/CIRC server configuration in the server.xml was also verified. And further Apache Web Server, Tomcat, and Oracle database connection were validated, and that OPACSVR/CIRCSVR is operational.
> (Duong, C. 2009 Mar 26. Focus Point of Open Platform, Voyager Web Services. at: Developers Meet Developers March 2009; [cited 2010 Nov 17]. Note: Some presentations

from this meeting are available at: http://www.exlibrisgroup.org/display/presentations/Developer+Meets+Developer+March+2009)

[Ho, Birong. "The Integrated Library System's APIs, an Open-source Web 2.0 Catalog, and University Computing Live Happily Ever After." Code4lib Journal 12, 2010-12-21 http://journal.code4lib.org/articles/4165]

Alma and implementation status at Princeton

Princeton University Libraries has been working with Ex Libris as a development partner for Alma, the company's new generation library management platform. As a development partner, Princeton has been able to work with each of the incremental releases of the software, testing functionality, with opportunities for input on the development of the product. Other development partners include Purdue University, Boston College, and KU Leuven. Boston College became the first library to place Alma into production use on July 2, 2012.

Princeton was represented on the Expert Advisory Group to advise Ex Libris on the Community Catalog of Alma by Katharine Treptow Farrell, Head, Order Division, and Assistant University Librarian for Technical Services.

The following excerpt provides an overview of Alma:

> Alma aims to strike a balance between providing automation services through shared resources and those that are managed locally. Its conceptual approach includes multiple layers, or zones:
>
> - The Community Zone provides a layer of broadly shared services. The community zone will include a large metadata store of bibliographic that Ex Libris will initially provide accessible by all Alma libraries. This store will also grow through contributions made by libraries as they create new records. Records in the community zone can be enhanced as needed by libraries using Alma.
> - The Library Zone is a metadata store specific to a single organization implementing Alma. Libraries may have various circumstances which might lead them to maintain some of their metadata in a private way.
>
> A library's holdings are managed through another data store called the Inventory, which can associate with descriptive data from either the Community or Library zones. As libraries add new items to their collections, they can tag onto existing records in the Community Zone, create new records in either zone as needed, and can retrieve records from external bibliographic services as needed, which can then be contributed to either zone, depending on concerns like record use restrictions. Ex Libris intends the metadata in the Community Zone to be available openly.
>
> The global services included with Alma will allow libraries to leverage widely shared resources. The bibliographic records available in the Community zone will facilitate efficient workflows when adding new materials to their collections. One of the broad goals of Alma is to allow libraries to execute common tasks in the most efficient ways so that they can concentrate more on the unique activities of more interest and strategic value to their organizations.
>
> The Metadata Management System provides tools and services related to metadata. This component will include editors and other tools that are involved in the creation and maintenance of

metadata. ALMA will take a format agnostic approach to metadata. Since it provides services for all types of materials, it will not be tied to MARC, but will be able to provide editing and validation tools for all applicable formats. The Metadata Management System will also support management of licensed electronic resources, relying on the knowledgebase currently associated with is SFX link server.

Ex Libris will offer Alma through software-as-a-service, hosted in a cloud infrastructure. During the development phase the system is hosted in Amazon's Elastic Compute Cloud. The company has not yet determined whether the production product will be delivered through the Amazon EC2 platform or some other competing infrastructure. The cloud infrastructure that Ex Libris ultimately selects will be based on cost, performance, and service-level agreements. While it's important to understand that Alma has been designed for deployment through the cloud, the specific provider ultimately selected will not necessarily be apparent to the libraries that use the product. [Ex Libris has since announced they will be using Equinix as its cloud infrastructure provider.]

Ex Libris has indicated that it intends to support local installations of Alma. Large consortia, for example, may prefer to implement and manage their own instance of URM. While the product has been designed for delivery through SaaS, some libraries may have specific needs that require a local installation.

The staff interfaces for Alma will be entirely Web-based. No local software will be required for library personnel to operate Alma and it is expected to function will all major Web browsers.

[Breeding, Marshall. "Ex Libris makes progress in delivering URM." *Smart Libraries Newsletter*, Vol 31 No 1. January 2011. Alma substituted in text for URM to reflect current branding.]

Although Princeton University has been working with Ex Libris for Alma, it has not set a specific date by which it plans to implement Alma as its production environment. The library plans to wait until it considers the software finished. It has developed a list of functionality that must be in place before going live with the software.

Informally, a transition to Alma might take place in December of 2012 or in July 2013 depending on progress made by those dates.

Depending on the dates of implementation of the technology for the next phase ReCAP, there should be a short-term plan to build interfaces into the current Voyager implementation as well as longer-term integration with Alma.

Ex Libris emphasizes that Alma will expose a robust set of APIs in addition to the functionality delivered through its staff and public interfaces. These APIs seem very well suited to the kinds of interoperability that will be needed to connect systems together more seamlessly in the technology to support the next phase of ReCAP.

Princeton University staff report, however, that so far in their experience with Alma as development partners that they have not had the opportunity to work with its APIs. The focus of effort has been on the delivered functionality. Any work with the APIs would have helped to reinforce their use by the

larger ReCAP project. The availability of the APIs will need to be further verified, including which APIs will be exposed in what releases of Alma.

Discovery Strategy

Princeton University currently has a variety of end-user interface and discovery tools, including the native WebVoyage online catalog delivered with Voyager, Primo, and Summon.

WebVoyage continues to be offered and is used by those with advanced research needs. It offers additional search options, including the ability to browse the collection based on authority headings or call numbers. Princeton does not plan to withdraw this interface until similar capabilities are available in its other discovery services.

Princeton has also implemented the Primo discovery product from Ex Libris. The local Primo index is currently populated primary with the records from Voyager. Records exported from Voyager are processed through a process, called a "pipe" in Ex Libris terminology, that normalizes the data as it is indexed. The Primo index is updated three times a day with added, modified, and deleted items.

Primo displays the current status of materials through a real-time query of the underlying management system, which is Voyager in the case of Princeton. This real-time availability is accomplished in one of two ways. One alternative uses a Web service built into Voyager. This service does not work well for some items, such as those that have hundreds of associated items. Princeton has also developed a custom program that retrieves the MARC, holdings, and item records associated with a title which works reliably even for those with extensive holdings data.

Patrons can choose to use Voyager catalog or Primo from the library's main Web page. No specific statistics have been provided, but the perceptions are that use of Primo is increasing relative to WebVoyage.

Princeton has not implemented Primo Central from Ex Libris which provides a deep index of articles, book chapters, and other resources, profiled according to the library's holdings. At the time that the library was ready to provide this kind of service, it determined that Summon provided more comprehensive coverage of its subscriptions to electronic resources.

Summon from Serials Solutions is offered as a separate discovery tool for articles. Materials from Princeton's Voyager ILS are not loaded into Summon, as do libraries that intend to use it as their comprehensive discovery service. Princeton currently positions Summon primarily as discovery tool specialized for e-journal articles and other electronic materials. Since it does not include books and other local materials, it does not need to interface directly with Voyager.

Princeton has developed an online form which has been integrated into both WebVoyage and Primo for ReCAP requests. The form is offered for materials coded with ReCAP locations. The form collects patron and item information, and lists the valid pick-up locations available. The form does not make direct updates within Voyager to the item status or patron record.

Data created the ReCAP request form is formatted according to the specifications developed by partner institutions and transferred by FTP to the GFA system three times per day. In the past each request was transmitted instantly, but was changed to scheduled batch transfers since some requests were lost. The scheduling of the batch transfer is scheduled to be available to ReCAP staff as they generate the pick lists each day.

Princeton does not have workflow management tools in place to manage the handling of ReCAP requests and returns. When a request is made that cannot be fulfilled by ReCAP, an error report is generated and transmitted to the library. These error reports are reviewed by library staff. In some cases the problem relates to requesting items that are not actually available to that patron or for the requested pick-up location, or that the item number does not exist in ReCAP, or that it has since been requested by another patron. The batch-oriented nature of the current processes lead to many opportunities for items to appear as available when they are not. Having more dynamic status information as materials flow in and out of ReCAP would address this issue.

One of the most common errors that Princeton reports in ReCAP requests involves invalid pick-up location. Any new technology tools created should fully validate all parameters of the request.

The request process does not automatically inject the request into Voyager. The request does not appear in the patron's account and the item status remains as available. As items are received from ReCAP, a staff member will manually place the item on hold for the patron account based on information on the printed ReCAP slip. Once placed on hold, a notice is sent to the patron. Items are then charged through normal circulation procedures.

The patron authentication environment at Princeton University primarily relies on barcode number and pin (last name) login credentials held in Voyager. This sign-in process is used by patrons as they access their account either through WebVoyage or Primo. The campus maintains a LDAP-based authentication service which is used by the library for Borrow Direct requests. Princeton has not yet implemented a single sign-on environment that spans all the services available to students, faculty, or staff. A Shibboleth service is available but has not been implemented extensively. The HathiTrust is one example of a resource accessed via Shibboleth.

Final decisions have not been made regarding how Alma might fit into the campus-wide authentication environment. Initially the implementation would be based on similar barcode and pin authorization as is now done in Voyager.

Summon does not require authentication for search. Access to items linked to in Summon depends on IP authentication, based on physical presence on campus, remote VPN access, or through a proxy service. Display of full-text subscribed resources requires authentication.

Princeton uses SFX and its associated knowledge base as its primary link resolver. The operation of Summon requires using the 360 Core from Serials Solutions to select the resources to be activated in the index. This arrangement requires keeping the 360 Core knowledge base (KnowledgeWorks) in synch with the SFX Global KnowledgeBase to the fullest extent possible. Princeton staff have learned that

packages and portfolios may often be organized differently in these two products, making it difficult exactly match all resources.

Princeton selected Summon because it determined that it more fully represented its electronic holdings. Primo Central has expanded its coverage since that decision was made. At some point Princeton may revisit its discovery options, but for the time being they plan to maintain this hybrid system.

New capabilities for the ReCAP will need to be created primarily with the Primo interface in mind, but should also take into consideration the expected ongoing use of WebVoyage and Summon.

Primo includes a capability Ex Libris calls "Deep Search" formerly referred to as "Third Node" that integrates access to resources that have compatible index structures to achieve immediate search capabilities without having to ingest the individual items in to the local instance of Primo. This capability works only for resources. Additional investigation will need to be performed to determine whether this capability can be anticipated to work well for blending an index of the shared ReCAP into Princeton's Primo instance.

One alternative to integrating the ReCAP Shared Collection into discovery layers such as Primo would be to export records form the ILS of the partner institutions.

The ETH-Bibliothek in Zurich has created a "Primo CMS Colr Websearch Adapter." That enables a Primo installation to connect directly to any SOLR server and directly integrate content into Primo.
http://www.exlibrisgroup.org/display/PrimoCC/Primo+CMS+Solr+Websearch+Adaptor

Columbia University Libraries

Columbia has approximately 3.7 million items currently stored in ReCAP.

Automation Strategy

Columbia University, like Princeton, uses the Voyager ILS as its primary automation environment. These two university library systems have entirely separate instances of Voyager. Columbia has been using Voyager since 2002 when it migrated from NOTIS.

Columbia University was a participant in the Kuali OLE project (kuali.org/ole) to create on open source enterprise-ready library management environment for research libraries during its initial 1-year planning phase (2008-2009). It did not join the subsequent 2-year project to build and implement the software.

2CUL

Columbia University and Cornell University have engaged in a partnership called 2CUL to collaborate which will include a "broad integration of resources collections, services, and expertise" (2cul.org). One aspect of this collaboration will include the implementation of a shared library management environment, with Alma as the primary system under consideration. 2CUL brings Columbia and Cornell into a close partnership while still maintaining their separate organizational structure and identities. The partnership will result in collaborative collection development, shared technical processing, and reciprocal access to collections by their respective patrons. 2CUL has received support from the Andrew W. Mellon Foundation.

Columbia's involvement with 2CUL has implications for the ReCAP project. Should 2CUL result in an automation environment shared between Columbia and Cornell, details will need to be worked out relative to Columbia's relationships to ReCAP that may not apply to Cornell.

Given that the commitments for Columbia to migrate from Voyager to Alma are not firm, it should be assumed that short term development should target the existing Voyager and Blacklight environment.

Authentication Environment

Columbia University maintains a central authentication service. The University has created an Apache module called mod_auth_pamacea to implement PAM (pluggable authentication modules) authentication. The Columbia University Libraries use PAM in conjunction with the EZ Proxy from OCLC to perform authentication for restricted resources and services.

The Arthur W. Diamond Law Library of Columbia University uses a Millennium system from Innovative Interfaces.

Discovery Environment

Columbia University has developed a discovery interface based on Blacklight. Blacklight is an open source discovery interface based on the Ruby on Rails development environment and the Apache SOLR search technology (projectblacklight.org). Original development work for Blacklight was initiated at the University of Virginia and it has been adopted by other major institutions such as Stanford University, the University of Wisconsin, Johns Hopkins University, North Carolina State University, Indiana University, the University of Hull in the UK, and by WGBH in Boston.

The beta version of Columbia's CLIO Blacklight catalog has been made public since March 2012: http://cliobeta.columbia.edu/

CLIO Beta follows a segmented approach to delivering search results. Parallel listings are provided for the 7.7 million volumes managed in Voyager and for those that come from Summon, the Academic Commons and other sources.

Columbia uses the 360 Suite of products from Serials Solutions for management and access to electronic resources.

Columbia University began subscribing to the Summon discovery service from Serials Solutions in 2010. It uses the company's 360 Core to profile its e-resource holdings and 360 Link as its OpenURL link resolver.

Resource Sharing Programs and Services

Columbia participates in **Borrow Direct**, which uses interlibrary loan management software from Relais International and the MasterKey federated search technology Index Data as the basis of an expedited book request and delivery service. Other participants in Borrow Direct include: Brown, Columbia, Cornell, Dartmouth, University of Pennsylvania, Princeton and Yale. Borrow Direct has been in operation since 1999. It was originally implemented using the URSA system from SirsiDynix, which has since been discontinued. The Borrow Direct system communicates with the ILS of each participating

institution using NCIP (NISO Circulation Interchange Protocol). Information regarding Columbia's policies regarding Borrow Direct: http://www.columbia.edu/cu/lweb/requestit/borrow_direct.html.

Columbia University offers a **Faculty Document Delivery Service** that allows faculty affiliated with the Morningside Heights Campus, Barnard College, and Union Theological Seminar to request articles held in the Butler Library to be scanned and delivered electronically.

Columbia participates in **MaRLI**, which gives selected faculty and graduate students access to collections of NYPL. This arrangement is done within the existing borrowing environments of the three institutions and does not involve any additional supporting technology.

Columbia also participates in the **RapidILL** service (rapidill.org) for resource sharing of articles. RapidILL was developed and is managed by Colorado State University. RapidILL was initially implemented at Columbia University as a centralized service, but it is being phased onto a more decentralized operation.

The Avery Architectural and Fine Arts Library has an electronic document delivery service that uses the ILLiad interlibrary loan management software from Atlas Systems.

Columbia University became a partner in the Google Books mass digitization project (announced December 2007) and joined the HathiTrust shared digital repository in December 2009.

ReCAP requests

Request from Columbia Patrons for recap materials take place primarily through an unmediated process using forms presented through the online catalog and discovery interfaces. Consistent with the processes implemented by the other ReCAP institutions, a Web form in the online catalog or discovery interface produces text file with the data representing the item and patron. Request data are aggregated into a text file that is then transmitted at specified times to the ReCAP GFA system.

A form for patrons to request off-site materials at ReCAP has been implemented in both the native Voyager catalog and in Blacklight CLIO. Request services have also been implemented for interlibrary loan and Borrow Direct.

Figure 7CLIO Beta record display with ReCAP request

The processing of ReCAP materials to fulfill patron requests for materials in ReCAP is handled by the various circulation service points among the Columbia University Libraries. These service points serve as pick-up locations for ReCAP materials. Interlibrary loan, in contrast, is operated more as a centralized service.

ReCAP Issues identified by Columbia

Columbia University is interested in a more seamless approach for borrowing ReCAP materials from the other partner institutions. The current approach works more like ILL. Current methods and procedures are focused on retrieving their own material from the facility, with no easy way for patrons to discover and request materials from the other two institutions.

Columbia personnel characterized the current ReCAP procedures as delicate and that they don't always work well. The batch processes involved for transmitting ReCAP requests have problems. The current processes do not accurately reflect what items are actually available. It is important that the new system be able to handle multiple simultaneous requests, which is not possible in the current batch process.

As general background, Princeton personnel report that their faculty members continue to be uneasy about ReCAP. Faculty lament that they cannot browse the stacks for all library materials. ReCAP must have excellent performance in its service to help library customers accept this arrangement. Service expectations will only increase over time. Princeton wants to promote ReCAP not only as increasing the amount of collection materials available to their patrons, but also as an improvement in service.

Columbia expects the new system to be very intelligent about the way that it exposes the shared ReCAP materials through its discovery interfaces.

Part II. Technology Implementation Options

This section provides an overview of concepts and issues relevant to the next phase of the technology support for ReCAP. It outlines some of the general requirements needed to fulfill the project goals related to management and access of a shared ReCAP collection.

Discovery of ReCAP Materials

In order to support the concept of a shared ReCAP collection with the ability for patrons to directly request materials, enhancements will need to be made to the discovery interfaces and online catalogs maintained by the three consortium members.

Key facts and assumptions:

- NYPL, Columbia, and Princeton University each use different automation systems and discovery services.
- There is no expectation that the three institutions would converge on a single discovery environment.
- End-user access to the ReCAP materials should not be offered through a separate interface.
- Materials available to their patrons through the shared ReCAP collection should be available through each institutions existing discovery interface, or through new interfaces that may be implemented in the future.

The Library Archival System developed by Generation Fifth Applications operates as an efficient inventory management system and does not include descriptive bibliographic data. One of the key considerations for the new ReCAP technology environment will be to not disrupt the existing operational efficiencies of LAS as it is used to manage the ReCAP facility.

Bibliographic data related to ReCAP materials resides only in the automation systems of the three partner institutions. The barcode number on an item provides the key linkage between bibliographic data and the ReCAP inventory.

Status and availability of materials in ReCAP is maintained in each ILS. Holding codes within each partner ILS are changed to ReCAP locations. ReCAP Customer Codes may or may not be stored in the ILS.

Status and availability of materials is maintained in the ReCAP GFA System. No systematic synchronization is performed between item status in GFA and the ILS of the three participants. In the current environment, the status of an item remains as available between the time it was requested and when it is pulled and shipped to the library.

Participants have indicated that they would prefer not to have to load the MARC records representing the holdings of the other ReCAP institutions into their local automation systems. Cross loading MARC records across partner ILS implementations would introduce significant costs and support burdens. In

some cases current system capacity or license thresholds would be exceeded. While cross loading MARC records is presented as a theoretical option, it is not a desirable alternative.

The ILS of the owning institution does not currently maintain real-time availability of an items throughout the lifecycle of the request, fulfillment and return process. Items continue to show as available between the time a patron request is made and when it is processed in the library for patron pick-up. There are significant intervals when even the owning library's ILS shows items as available when they are not.

A discovery infrastructure will be created to provide improved patron access to the shared ReCAP materials. At a minimum, this discovery component will be populated with bibliographic and administrative metadata for the items identified as belonging to the shared ReCAP collection.

Option: Should the discovery environment also be populated with a broader selection of ReCAP records? Index views could be defined to scope search to the Shared ReCAP collection for patron search and other views defined for library personnel to support collection management activities. The design of the new infrastructure should accommodate multiple scenarios for what records are harvested and indexed.

Is there a need to filter inclusion of ReCAP materials by institution? Are there examples where materials should display in the ReCAP index of the home institution and not the partner institutions?

Should there be a status flag in the patron records for MaRLI which expands a patron's eligibility for ReCAP materials? What other kinds of patron scenarios need to be accommodated in the ReCAP discovery and delivery environment?

Expected functionality for ReCAP Discovery

Discovery of ReCAP materials should not require users to go to a separate service. New functionality will be created that allows patrons to see the shared ReCAP materials among search results and to place requests for them. Authentication of patrons will rely on mechanisms currently in place.

Discovery of ReCAP materials should be integrated into each the discovery services or online catalogs in general use for the participating institutions.

- Issue: The creation of an index for ReCAP materials could likely be integrated into discovery layer products such as Blacklight, Primo, or BiblioCommons designed for segmented index architectures. Integrating this capability into online catalogs will be much more of a challenge without having to resort to loading the MARC records for the shared ReCAP collection into the respective ILS instances.

ReCAP materials for each library's home institution will already be present in their ILS which directly drives its online catalog and is used to populate its discovery layer.

The key extension for the next-phase ReCAP involves the inclusion of materials from the other two institutions.

Search results should appropriately label materials as “off-site” and maintain internally the GFA Customer Codes or other factors needed to calculate eligibility for request for each request context and valid pick-up locations.

New system should provide a service that provides accurate data in real time regarding whether an item is available for request. This real-time availability status service would be consumed by the discovery or online catalogs of the partner institutions and may be useful to other business processes.

ReCAP materials subject for discovery will include all types of materials: monographs, serials, media, etc.

The infrastructure should be flexible regarding how items are defined as belonging to the ReCAP Shared Collection. Possible factors for inclusion criteria might include GFA Customer Codes, ILS holding locations, item types, date ranges, or designated individual items.

At this phase recap discovery will include journals at the title and issue level. It will not directly include article-level data.

Any article-level discovery products used by the partner institutions should be able to appropriate link to items in the Shared ReCAP collection should be handled through OpenURL link resolvers or other mechanisms. Resolver menus would offer electronic document delivery options for these materials.

Within each discovery environment, ReCAP materials should be capable of being interfiled among search result listings or offered as a separate result set, depending on library preference or by the user through facet selection.

Through what interfaces is discovery expected for the Shared ReCAP collection:

- NYPL WebPAC Pro online catalog
- NYPL BiblioCommons discovery service
- Princeton University Main Catalog (WebVoyage)
- Princeton University SearchIt@PUL (Primo)
- Columbia University Classic Clio Catalog (WebVoyage)
- Columbia University CLIObeta (Blacklight)

ReCAP results will not be included in:

- Princeton University Articles + (Summon core index only)

Technical implementation options

Cross-loading Records in each ILS

One approach that would be the most technically feasible but include the most overhead would involve loading the records from the partner institutions into each other’s ILS implementations:

- Load Princeton and Columbia sharable ReCAP records into the Millennium ILS of NYPL
- Load NYPL and Columbia sharable ReCAP records into the Voyager ILS of Princeton

- Load NYPL and Princeton sharable ReCAP records into the Voyager ILS of Columbia

This approach has disadvantages:

- It would introduce millions of records into each automation environment which would need to be maintained.
- It could greatly increase ILS licensing costs since fees include as a factor the number of bibliographic records managed
- It would greatly increase technical services processing needed to maintain authority control and other database quality assurance issues
- It would introduce significant work to reconcile serials holdings represented in the ILS based on materials in the other ReCAP institutions

Support of this model would include several procedural steps and technology components:

- Selective extraction of records from each of the three ILS implementations
    - Extraction would include full MARC records with holding and status information encoded on 9xx fields
    - Current status data will be included on the extraction and data loading routines, but will be definitively displayed based on the new real-time availability service.
    - Selection of records to be extracted would be based on criteria defined by the ReCAP participants. Factors for the criteria might include: ReCAP Customer Code, ILS holding location, Item Type, Date Range, or individually designated items.
        - Options for selection would include whether to include all ReCAP customer codes or just those designated as sharable by other partners
    - Extraction would be performed though standard procedures as described in the DLF Integrated Library System – Discovery Protocol (ILS-DI), which generally corresponded to the Open Archives Initiative Protocol for Metadata Harvesting (OAI-PMH)
    - Extracted records would be held in a common file system for later processing
    - Incremental extracts would be performed at least daily for new items transferred to ReCAP, for withdrawals, or of other changes
    - Supporting technology components might include:
        - OAI Toolkit from the eXtensible Catalog project
          "The Metadata Services Toolkit enables the XC user interface to present FRBRized, faceted navigation across a range of library resources. The toolkit aggregates metadata from various silos, normalizes (cleans-up) metadata of varying levels of quality, and transform MARC and DC metadata into a consistent format for use in the discovery layer."
          http://www.extensiblecatalog.org/
    - Extracted records would also need to be extended with date derived from the GFA Library Archival System.
        - Based on the barcode number as the unique match point corresponding records would be retrieved from GFA

        - Customer Code (designates owing library and collection and used to determine policies on eligibility for request, loan periods, etc.)
        - Status information – availability
        - These data points would be encoded on the 9xx fields since they would be needed to support a preliminary status display in discovery systems and request eligibility. [Are ReCAP customer codes already present in ILS bibliographic records? Can they be definitively calculated? It would simplify the process if retrieval of GFA records could be eliminated.]
    - Record Extraction routines will also be needed for other implementation scenarios
- Loading of MARC records into ILS systems
    - Records already exist for ReCAP holdings of each institution
    - Records would be loaded into each of the three ILS implementations of the MARC record sets from the other two institutions
    - The process would use the standard record loaders that parse the 9xx fields to create holdings and item records
    - Depending on the cataloging standards of each of the institutions, records may need to be processed to validate headings or other tasks needed to preserve the integrity of authority control
    - Possible processing required to deal with serial records and holdings
    - Processing should detect duplicates and carry out selected processing options:
        - Suppress duplicate records and add additional copy statements to existing bibliographic records
        - Reject and issue error report
        - Load duplicate record and generate informational message
    - Incremental loads would be performed at least daily to synchronize ILS holdings with ReCAP inventory across participating institutions
- Enable Real-time availability
    - The creation of a real-time availability process will be needed by any of the implementation scenarios and is described below.

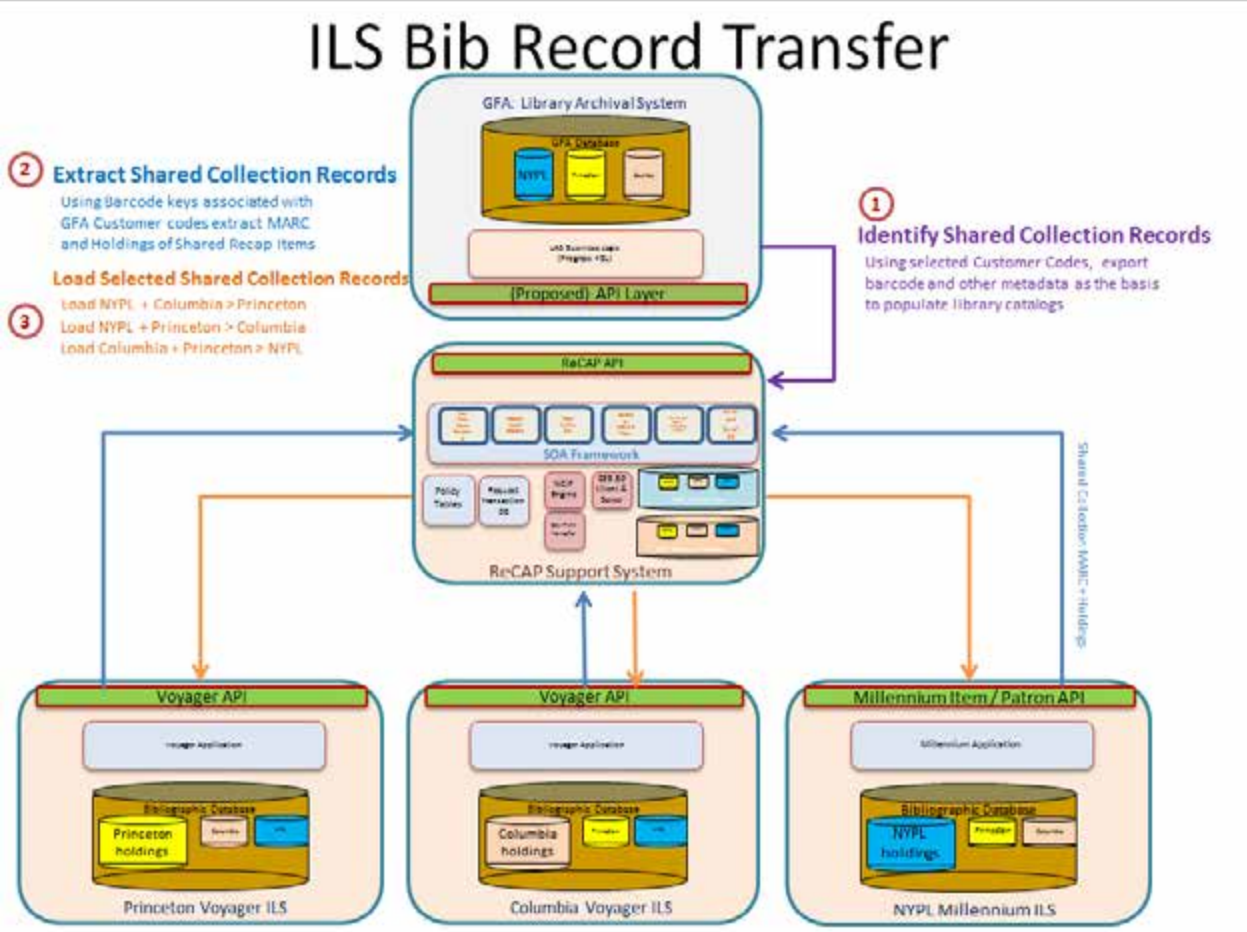

Figure 8 Diagram of ILS Cross institutional MARC record loading

All of the partner libraries[]

Creation of a modular ReCAP index or plug-in

This approach would involve the creation of an index that could be integrated into the relevant discovery services.

A discovery environment would be based on an index created from bibliographic records representing the Shared ReCAP collection. It would rely on a bibliographic database as a component of the ReCAP infrastructure created according basically the same requirements as those used for the bibliographic cross-loading option. The key difference is that the bibliographic database remains in the central ReCAP application and no additional records are loaded in the partner institution ILS implementations.

- Selective extraction of records from each of the three ILS implementations
    - Extraction would include full MARC records with holding and status information encoded on 9xx fields

- Current status data will be included on the extraction and data loading routines, but will be definitively displayed based on the new real-time availability service.
- Selection of records to be extracted would be based on criteria defined by the ReCAP participants. Factors for the criteria might include: ReCAP Customer Code, ILS holding location, Item Type, Date Range, or individually designated items.
    - Options for selection would include whether to include all ReCAP customer codes or just those designated as sharable by other partners
- Extraction would be performed though standard procedures as described in the DLF Integrated Library System – Discovery Protocol (ILS-DI), which generally corresponded to the Open Archives Initiative Protocol for Metadata Harvesting (OAI-PMH)
- Extracted records would be held in a common file system for later processing
- Incremental extracts would be performed at least daily for new items transferred to ReCAP, for withdrawals, or of other changes
- Supporting technology components might include:
    - OAI Toolkit from the eXtensible Catalog project
      "The Metadata Services Toolkit enables the XC user interface to present FRBRized, faceted navigation across a range of library resources. The toolkit aggregates metadata from various silos, normalizes (cleans-up) metadata of varying levels of quality, and transform MARC and DC metadata into a consistent format for use in the discovery layer."
      http://www.extensiblecatalog.org/
- Extracted records would also need to be extended with date derived from the GFA Library Archival System.
    - Based on the barcode number as the unique match point corresponding records would be retrieved from GFA
    - Customer Code (designates owing library and collection and used to determine policies on eligibility for request, loan periods, etc.)
    - Status information – availability
    - These data points would be encoded on the 9xx fields since they would be needed to support a preliminary status display in discovery systems and request eligibility. [Are ReCAP customer codes already present in ILS bibliographic records? Can they be definitively calculated? It would simplify the process if retrieval of GFA records could be eliminated.]
- Record Extraction routines will also be needed for other implementation scenarios

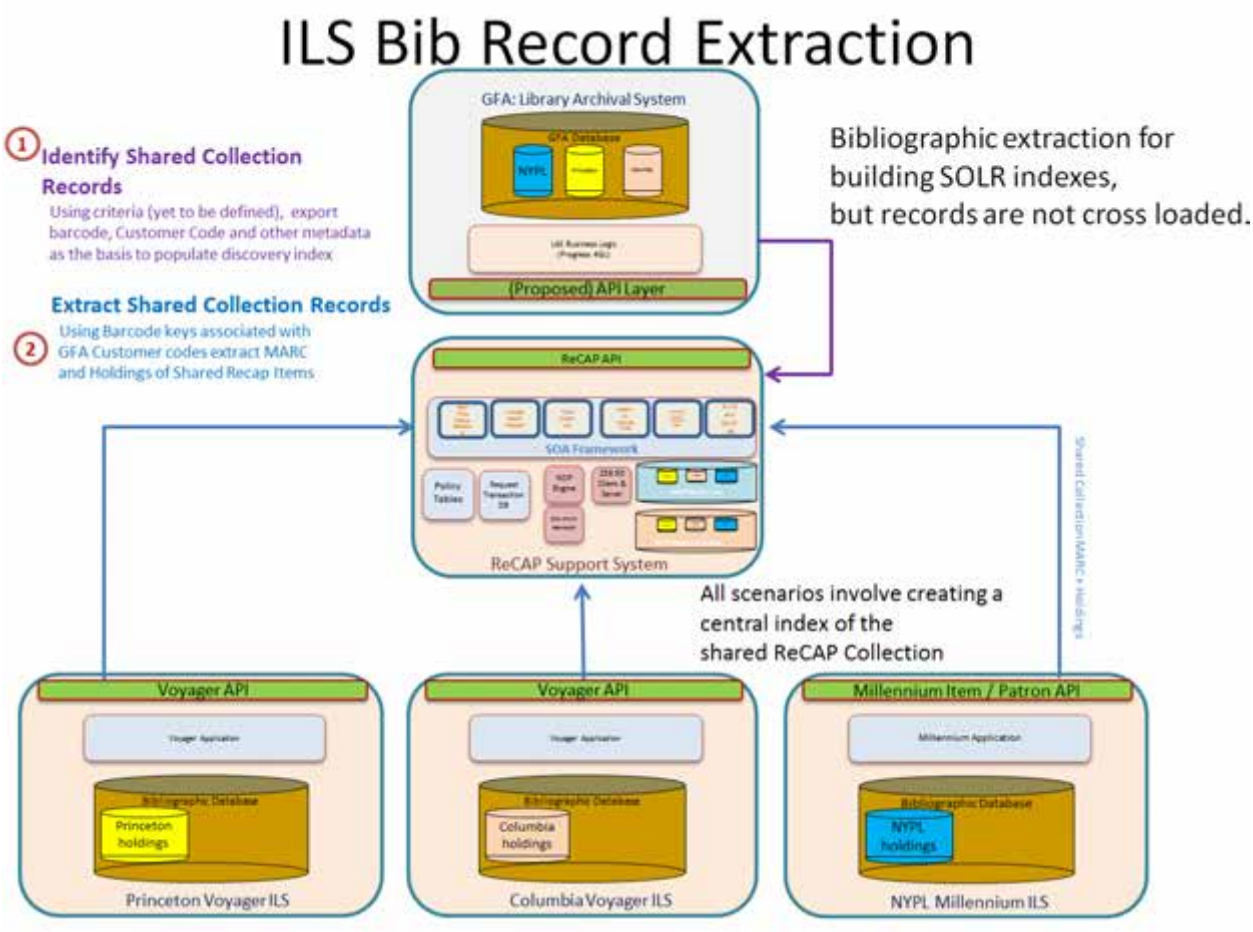

Figure 9 Bibliographic Record extraction for ReCAP Discovery

Apache SOLR / SOLRMARC

The obvious candidate technology to support this approach would be the SOLR and Lucene search and retrieval applications available as open source under the governance of the Apache Foundation.

One implementation option would involve using the general infrastructure of one of the SOLR-based discovery environments to provide a platform for creating indexes and for metadata management tasks. Options include:

- Blacklight: open source discovery services based on Apache SOLR Lucene and built using the Ruby on Rails development environment
- VuFind: open source discovery service based on Apache SOLR Lucene and built using the PHP programming language

One instance of this platform would be maintained on behalf of all the ReCAP institutions. The platform would need to be sized to accommodate indexing processes and other support tasks at the scale of the combined ReCAP inventory.

The SOLR index would be created by indexing the bibliographic records harvested and synchronized from the partner ILS implementation. Data would staged in the MARC-based bibliographic database and SOLRMARC used to build and maintain the index accessed by discovery systems.

It is would also be possible to associate the bibliographic database with a Z39.50 or SRU service to allow integration of the Shared ReCAP Collection by systems not compatile with Apache SOLR or Lucene.

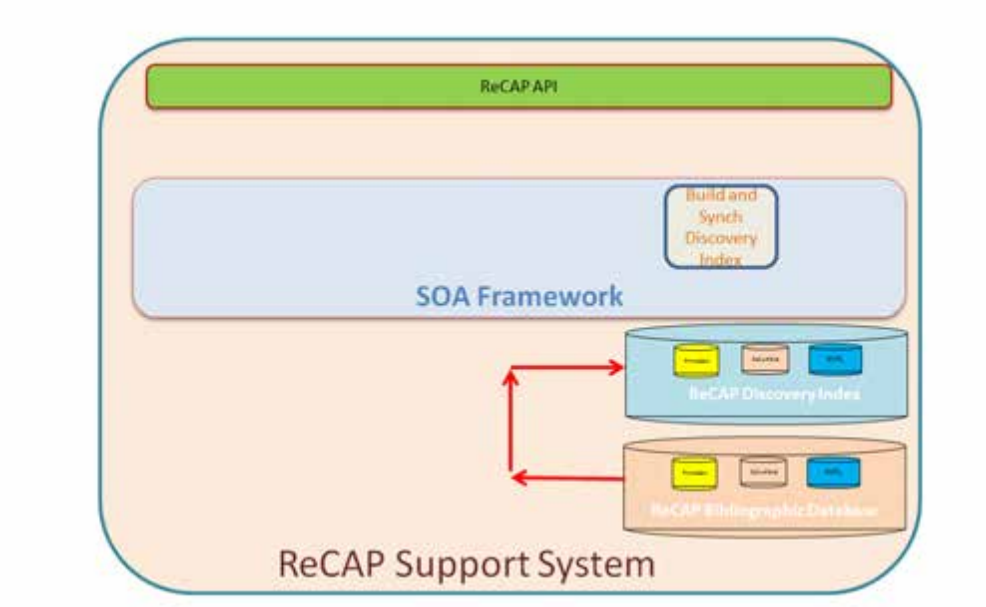

Figure 10 Figure 10 Creation of SOLR index derived from central MARC database

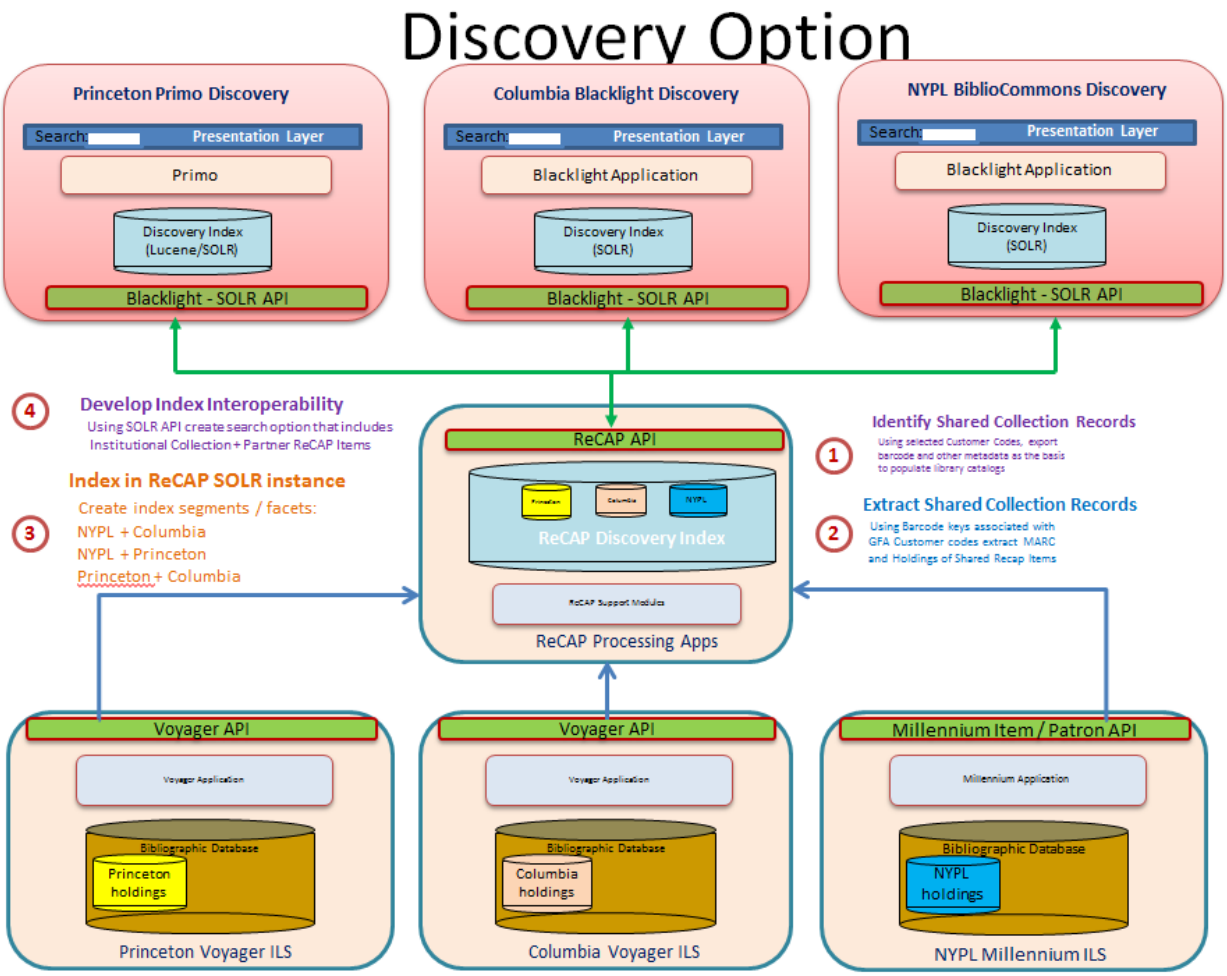

Figure 11 Model based on SOLR discovery platform

Additional metadata support would be provided through the tools developed through the eXtensible Catalog project:

- XC NCIP Toolkit
- XC Metadata Management Toolkit:
  The Metadata Services Toolkit enables the XC user interface to present FRBRized, faceted navigation across a range of library resources. The toolkit aggregates metadata from various silos, normalizes (cleans-up) metadata of varying levels of quality, and transform MARC and DC metadata into a consistent format for use in the discovery layer.
  This toolkit presents an opportunity for libraries to apply their expertise with creating and managing metadata in a variety of web applications.
  http://www.extensiblecatalog.org/

Data exchange and communications capabilities for the environment might include components such as:

- NCIP Toolkit Connectors for Voyager: http://www.extensiblecatalog.org/news/xc-ncip-toolkit-connectors-available-symphony-and-voyager-ils
  http://code.google.com/p/xcncip2toolkit/wiki/VoyagerInstallation
- Patron Account Information API
  http://gbv.github.com/paia/paia-868166f.html
- DAIA - Document Availability Information API
  http://www.gbv.de/wikis/cls/DAIA_-_Document_Availability_Information_API
  "The Document Availability Information API (DAIA) defines a data model with serializations in JSON and XML to encode information about the current availability of documents. This document defines the serialization formats DAIA/JSON and DAIA/XML and a HTTP query API to query DAIA information."

This approach will be technically challenging. It assumes an environment of discovery services capable of integrating with third-party indexes in addition to the one built-in to the service. Some of the considerations of this approach:

- Bibliographic records corresponding to each institution's sharable ReCAP records will be extracted from the respective ILS implementations as described above
- Records will be indexed in SOLR
- Facets or other mechanisms will be implemented to allow segmentation of the index according to the three owing institutions
- The environment will need to be capable of presenting indexing views based on selected institutional segments. The ReCAP SOLR index to be integrated should activate only the ReCAP items from the other two institutions and not duplicate indexing already present in their discovery environment for their own ReCAP items
- This approach will be challenging to include in the native online catalogs for the three institutions since they do not make use of discovery indexes but rather are generated directly out of the respective ILS implementations.
- All three discovery environments are based on SOLR:
  - NYPL BiblioCommons
  - Columbia's Blacklight-based CLIO beta
  - Princeton's Primo: based on Lucene. The Ex Libris "Deep Search" technology provides support for access to third party indexes in addition to records loaded into Primo through pipes. Primo Central, for example, integrates article-level materials with the local index using this Deep Search
- Additional investigation will need to be performed to ensure the feasibility of this index integration approach.
- If successful, multiple indexes within a discovery environment will work together to produce blended results ordered according to relevancy algorithms applied consistently among the sources. Index segments would include:

    - Primary discovery index derived from materials managed through the ILS implementation
    - Additional index representing ReCAP materials
    - Additional indexes representing article-level materials, such as Summon or Primo Central.
- Item display routines would need to interact with different business systems for real-time availability and status:
    - For items managed within the local ILS, the standard availability routines would apply. Each of the current discovery implementations includes dynamic interrogation of the ILS for circulation status and for patron My Account functionality
    - For items representing ReCAP items from partner institutions, item display (brief or full) would trigger a call to the Web service on the ReCAP GFA system to retrieve current status
    - The status of an item in GFA LAS would be definitive only if the status is updated at the initiation of a patron request. A two-factor availability status optionally could be implemented based on both LAS status and that from the relevant ILS.
- Request forms, described in more detail below, would be presented for ReCAP items. Signed-in users might trigger additional processing that validates eligibility for request.

Creation of ReCAP Requests

A standard request form should be designed that collects consistent information and that validates all data collected. Examples:

- Authenticate the patron through available services:
    - Through the proprietary sign-on mechanism provided through the online catalog or discovery interface
    - Through NCIP call into the ILS
    - Through a campus authentication service
- Reliable performance of the fulfillment requires that patrons be authenticated and that any required details are present in their patron record.
- Ensure that patron is eligible to make request
- Patron data either derived from or validated against patron record in home institution ILS
- Validate that the pick-up location is valid for that type of material and that type of user
- Validate that they item requested is actually available through dynamic status information in ILS and GFA
- The form should accommodate requests made directly by patrons through library discovery interfaces or online catalogs
- The form should also be designed to accommodate requests mediated through library staff members on behalf of patrons. This form would be used to enter requests made over the phone, send in by e-mail, etc.

- Staff-mediated requests would perform the same patron validation as when unmediated. The staff member might, for example, request the patron's patron ID number and the system would dynamically interrogate the ILS using NCIP or other applicable protocol, to retrieve the patron record and validate request conditions.

- ReCAP requests once validated by the patron would then be submitted to a workflow management application as described below.
- Automatically transmit the request to the GFA request queue
- Update GFA system dynamically upon successful processing of transaction to flag item as in use.
- Upon successful submission of the request, the a confirmation message will be displayed to the patron and optionally transmitted by e-mail
- The expected time of fulfillment should be displayed, based on when the request was submitted relative to the current schedule of ReCAP for pulling requests, courier schedules, library opening hours, and other factors.
- A transaction identifier assigned so that patrons and library or ReCAP personnel can check the status of the request throughout the fulfillment process.

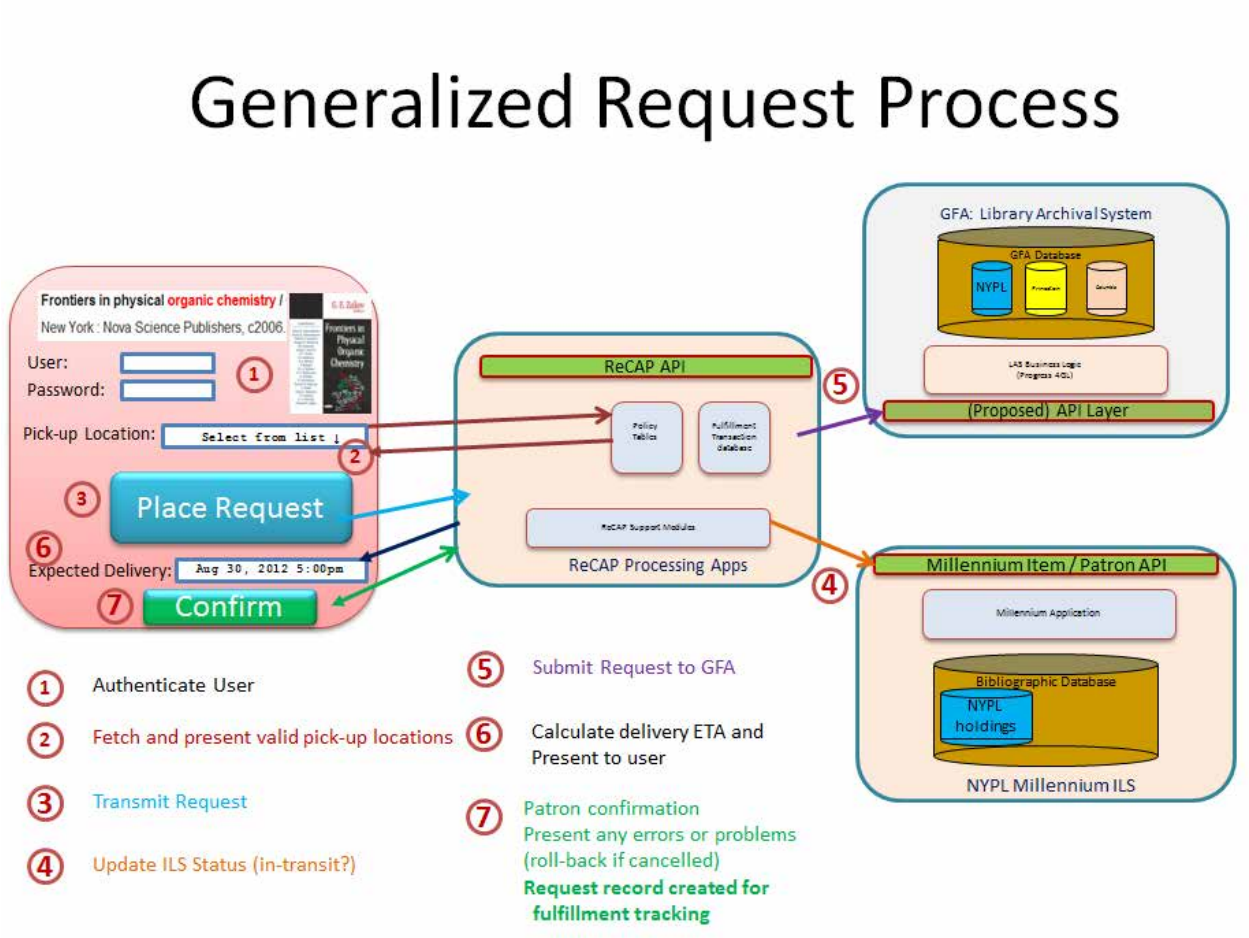

Figure 12 Proposed model for generalized ReCAP request processing

Real-time Availability Environment for ReCAP

One of the key requirements for the next generation of ReCAP technology support will involve the ability to display the status and availability of ReCAP item in real time. Some of the problems with the existing processes have been identified as related to the batch processes used for item requests and status updates. Also in order to extend the operational scenario to include requests across all participating institutions, support for real-time status information and the ability to dynamically update statuses based on current transactions will be essential.

Some of the options to implement a dynamic status environment might include:

Implement a layer of services that dynamically interacts with all four business applications:

- The GFA Library Archival System that manages the ReCAP inventory
- The Millennium ILS of NYPL or future automation platform
- The Voyager ILS of Princeton University or future automation platform
- The Voyager ILS of Columbia University or future automation platform

Add a services layer to the GFA Library Archival System that exposes a set of Application Programming Interfaces (APIs) preferably in the form of secure REST services or enveloped in SOAP.

- This service layer would need to be created by developers with proprietary access to the GFA ToolPro environment.
    - Such work would likely be accomplished as a sponsored development project with Generation Fifth Associates, but there might be other possible scenarios
    - Access to services would require authentication to ensure system integrity
- Examples of services needed to support ReCAP might include
    - Request item status: presentation of an item identifier (barcode number) would return the items current status and availability
    - Update status: presentation of an item identifier and a desired status (available, in process, in transit, charged to user, etc.) would return an error or success flag
- This service layer could also take the form of a NCIP responder. Selected NCIP operations could be implemented with the specific functions needed to maintain a dynamic status and availability environment. This approach could take advantage of the multiple open source implementations of NCIP available, including those offered through the eXtensible Catalog project.

Create functionality that can be embedded in each of the online catalogs or discovery layers that dynamically calculates the status and availability of ReCAP items that appear in search results

- The items will include location codes that cause display of something like “off-site”
- Items will also include access to the ReCAP Customer Codes to support calculation of eligibility for request and to validate pick-up locations
- A preliminary availability status might be presented based on statuses harvested from ReCAP records and updated through the current batch system

- Appearance of a ReCAP item during the presentation of a brief results list or on a full record display would trigger a call to the ReCAP server's availability service. This request might be implemented through AJAX (Asynchronous JavaScript and XML).
- This service could also be implemented

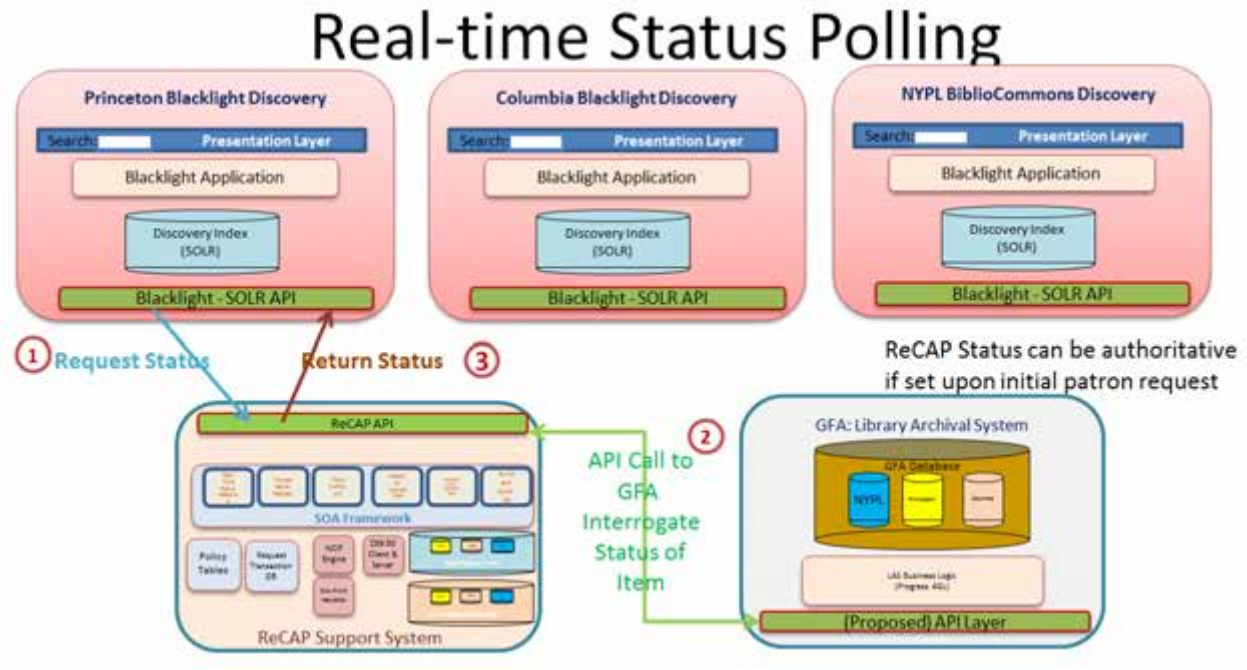

Figure 13 Single factor Model of Possible real-time status display

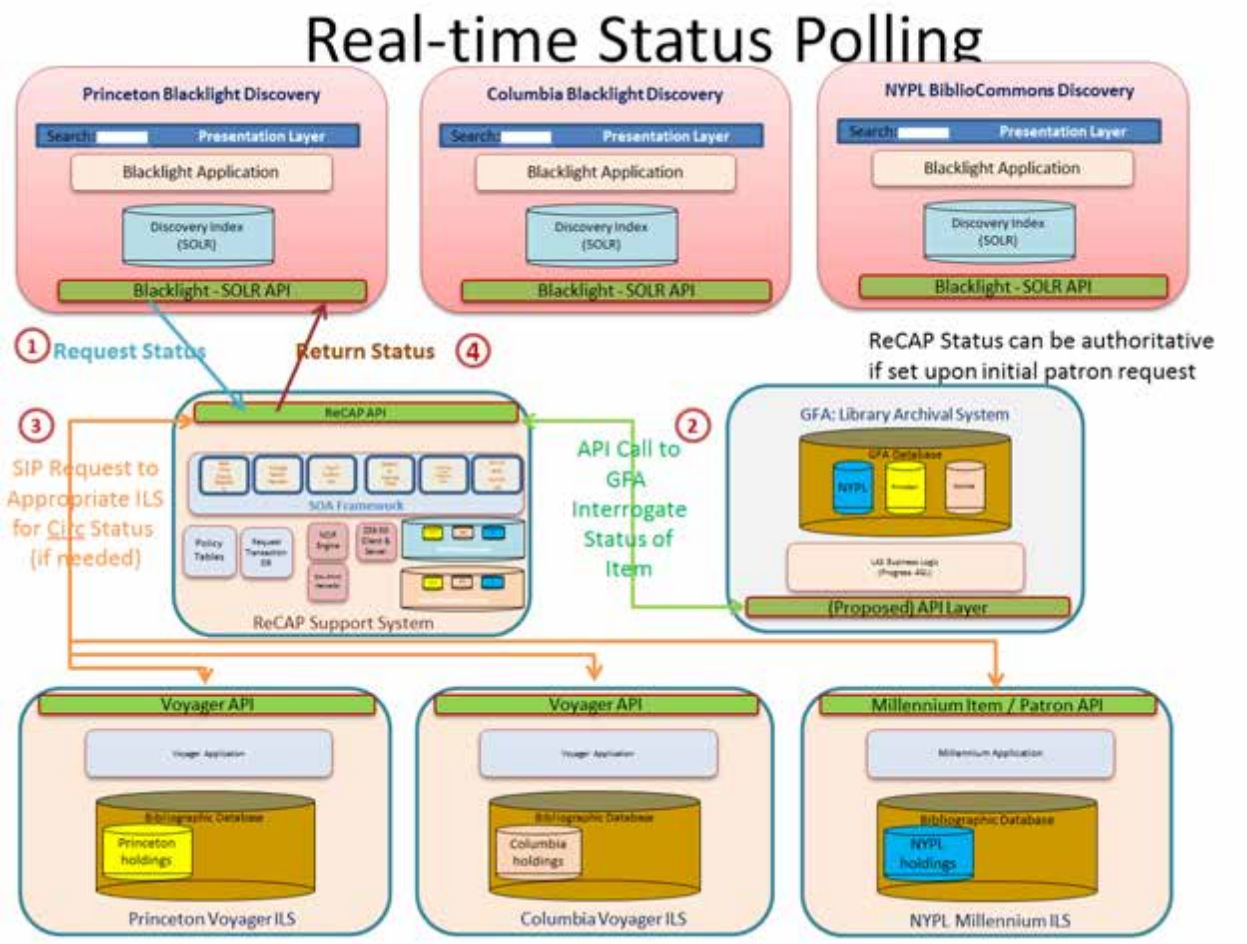

Figure 14Two-factor model for real-time status

Management of ReCAP Requests

The improvement of fulfillment services related to ReCAP requests could be accomplished through an workflow management system that automates, validates, and verifies all aspects of the transactions involving ReCAP materials that transpire in the gaps between the ReCAP GFA system and the respective ILS implementations. The workflow management system would also provide supporting infrastructure for GFA and the ILS implementations that might not be present natively.

The GFA system excels at ensuring that no items are lost within the ReCAP storage facility. It does not perform any tracking of items as they proceed through external business processes: transit to library, processing in library, check-out to patron, routing within library branches, and transit back to ReCAP. In order to gain better control of ReCAP materials, a complete end-to-end management system is required.

The Workflow Management Utility would be a central point for working with policies that apply to ReCAP materials. It would maintain a table of policies that includes a matrix of GFA Customer Codes, ILS patron types and eligibility, valid pickup locations, as well as parameters needed in support of electronic document delivery.

A fulfillment module could optionally be extended to other types of patron services beyond ReCAP requests. Such a an generalized utility would bring consistency and control to a wide range of patron requests, including holds and recalls, interlibrary loan, document delivery, in addition to ReCAP requests. This system would provide a consistent authentication and present a consistent user interface and take advantage of business logic that would present the most efficient options available to the patron for receiving a desired item, and would pre-validate all data to ensure successful fulfillment. Such an extension of capability would be outside the scope of the initial project, but it should be designed in ways to support workflows other than ReCAP requests.

Many features would be implemented as a Web service that could be addressed by any authorized discovery service of the participating institutions.

The Workflow Management Utility would maintain a transaction record for each request. Transactions would be opened when submitted through the ReCAP request form and would be closed when items are returned to the ReCAP storage modules.

Transactions would be retained in ways that preserve patron privacy but that support full reporting capabilities, including queries of pending transactions and retrospective performance statistics.

The Workflow Management Utility would implement or support a federated identity management system as needed to support authentication requests needed for ReCAP services.

Some of the features of the ReCAP Workflow Management Utility might include:

Business logic that determines the availability of a ReCAP item to a library user based on:

- Embedded ReCAP Customer Code of the item
- Patron institution and category

The utility would perform the processing and validation of the ReCAP request forms embedded on each of the online catalogs or discovery services. It would validate patron eligibility and calculate valid pick-up locations based on applicable parameters.

ReCAP requests should show in a patrons ILS account as would a hold or other request for library materials.

- This capability might be accomplished through a dynamic process, possibly implemented through an NCIP call that pushes a transaction to the ILS placing the item on call for the patron as part of the processing of the ReCAP request form.
- For non-local ReCAP items, this process might include the creation of a temporary item record in the ILS to support the expected subsequent circulation transaction.
- Return of items should likewise remove temporary item records

Patrons should be automatically notified when an item that they requested from ReCAP is available for pick up.

- This functionality would be accomplished by the local ILS implementation following the injection of the hold request made when the item was requested

When placing a ReCAP request a list of valid pick-up requests should be presented

- This feature would be accomplished by a dynamic call made by the ReCAP request form to the Workflow Management Module that would return the valid options for that patron category

ReCAP materials should be tracked through the entire fulfillment and return process. Patrons should receive something like a tracking ID that they can use to view the status of their request.

Materials should be scanned at each step of the fulfillment and return process, including:

- Pickup at the ReCAP facility
- Arrival at the library
- Arrival at the requested pick-up location
- Check-out to patron
- Check-in by patron
- Pick-up at library in transit to ReCAP
- Arrival at ReCAP
- Re-filed at recap into storage module

Workflow Management Utility should generate performance statistics to document percentage of transactions with on-time delivery

Workflow Management Utility should generate alerts of items that have not been delivered in expected time frames.

The Workflow Management Utility would include a Web form that patrons can use to inquire the status of a pending request. This feature would work much like the tracking systems used by parcel delivery services, displaying details of all previous steps in processing the request up to the current status.

A form would be available for library personnel to see the status of any given item or category of items.

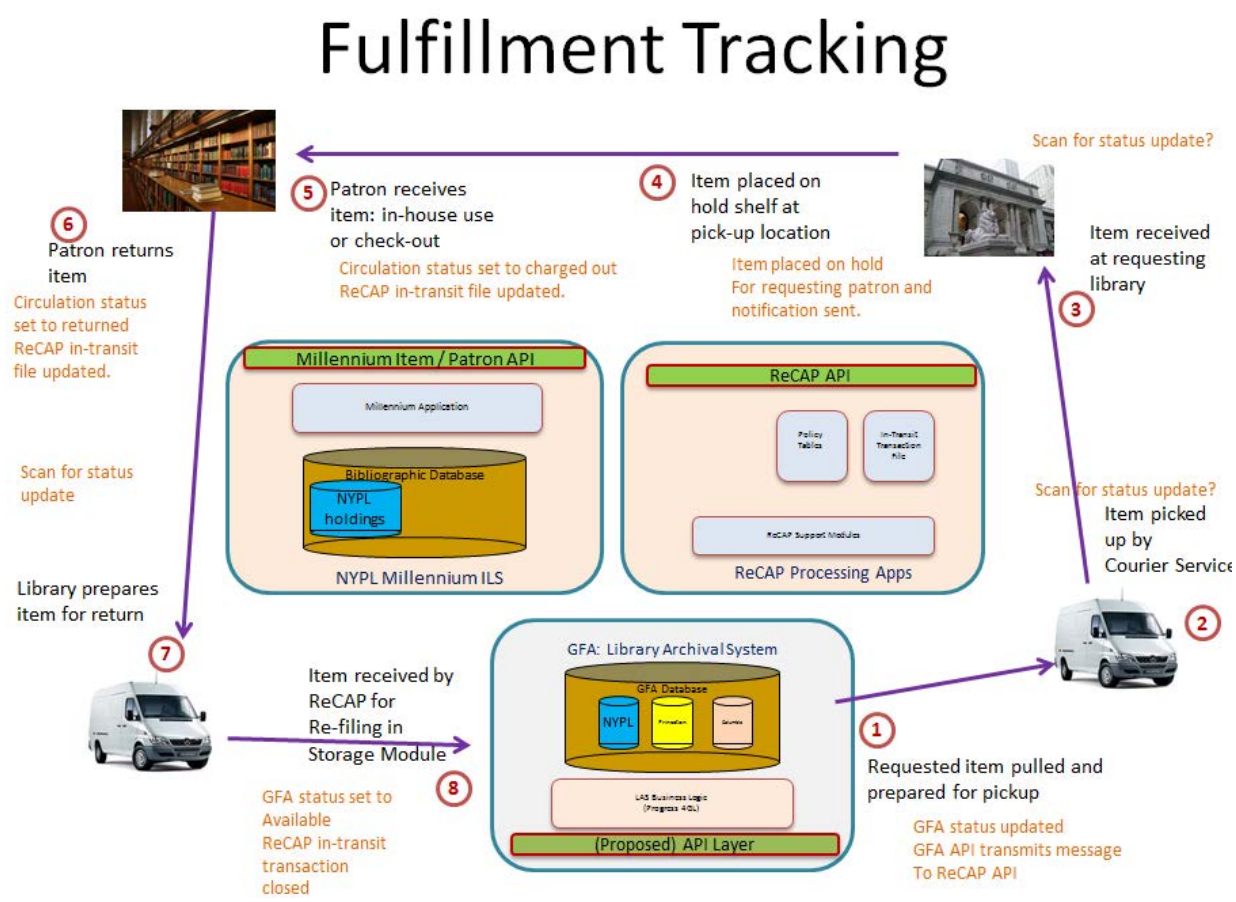

Figure 15 General flow of fulfillment transactions

Optional Functionality or future enhancements

Each of the three institutions participating in ReCAP also participated in the Google Books mass digitization project. They also are involved in the HathiTrust, a digital repository that holds the library copies of Google Book scans, in addition to other materials. One possible element of functionality would involve a query to HathiTrust to determine if the item could be fulfilled digitally. There would be limited circumstances where the copyright status of the item would allow such digital delivery.

Technical implementation

The new technology infrastructure might be called the Enhanced ReCAP Discovery and Delivery Service. This infrastructure will not be a standalone application, but rather a set of components assembled as middleware to provide services to the strategic business application and end-user discovery services in place in the participating institutions.

The proposed functionality described will be implemented in a set of tools that will be delivered through a service-oriented architecture. While some of the functionality might involve interfaces operated by library patrons or staff, much of the environment will consist of services exposed as part of the middleware application that will be consumed in system-to-system operations. These services will exploit existing APIs or protocol responders currently or potentially available in the related integrated

library systems and discovery services. The GFA Library Archival System lacks these APIs, and part of the scope of the development will be the creation of a service layer for this application.

Although the functionality has been designed specifically for ReCAP the intent is also to make the technical environment as generalizable as possible so that it could also be adapted to other organizations. To the extent possible the services supporting the business logic will be constructed using API's and protocols commonly used in the library and archives environment, including those defined in the Digital Library Federation sponsored ILS-DI initiative, the Mellon funded eXtensible Catalog project, NISO standards including OAI-PMH, NCIP, SIP, and SRU. APIs created specifically in support of the project should be created as RESTful Web services, authenticated and transported via SSL.

The design should be implemented in ways that provide the highest level of flexibility of possible. Many of the assumptions and business rules initially expressed to create these designs may be subject to reconsideration. To the extent possible, the business logic should be based on policy tables or parameters that can be adjusted as needed. Such flexibility will be needed for such fundamental constructs such as the definition of the ReCAP Shared Collection.

As a service-oriented application, the ReCAP Workflow Management Utility will be constructed as business logic and services that will operate on an existing middleware framework. Making use of one of the existing open source and commercial middleware environments already available would significantly reduce the effort of a development project rather than building all the low-level supporting services anew.

The selection of the low-level middleware framework to support the development of the ReCAP Workflow Management Utility requires further investigation. One possibility that stands out for consideration is Kuali Rice:

> The Kuali Rice software provides an enterprise class middleware suite of integrated products that allows for applications to be built in an agile fashion. This enables developers to react to end-user business requirements in an efficient and productive manner, so that they can produce high quality business applications.
>
> Kuali Rice is leveraged heavily by the Kuali applications but is also designed to be used in non-Kuali applications. Its services and framework pieces are designed in such a way to be applicable to multiple business domains.
>
> Rice is built with Service Oriented Architecture (SOA) concepts in mind. Specifically, end developers are able to build robust systems with common enterprise workflow functionality, customizable and configurable user interfaces with a clean and universal look and feel, and general notification features to allow for a consolidated list of work "action items." Additionally, there are a set of services in Rice that provide identity and access management capabilities and can be used to abstract away from underlying

institution-specific identity services. All of this adds up to a re-usable development framework that encourages a simplified approach to developing true business functionality as modular applications.

http://www.kuali.org/rice

The use of Kuali Rice for low-level infrastructure would provide a solid technical environment and would be consistent with any requirements for software developed under the current or follow-on grants to be made available as open source. The Kuali Foundation is closely aligned with the Andrew W. Mellon foundation, the key funding source for this ReCAP current planning effort. Kuali Rice is used by the Kuali OLE project as its low-level service bus and workflow support layer.

Other development frameworks could likewise provide the foundation of the ReCAP technical support infrastructure. IBM, for example, offers a SOA development framework called the IBM Integration Designer. http://www-01.ibm.com/software/integration/integration-designer/

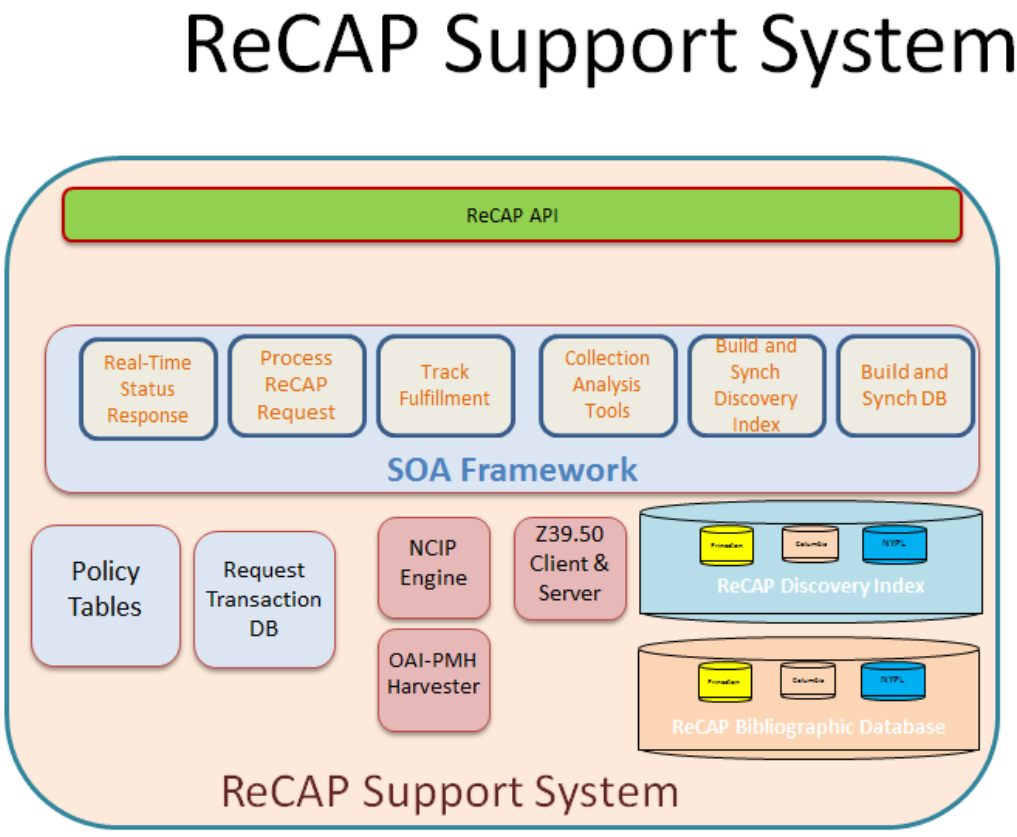

Figure 16 Simplified model of ReCAP support architecture

Support for Electronic Document Delivery Requests

The fulfillment of electronic document delivery requests would also benefit from the enhanced technical infrastructure. While no physical tracking will be needed, tracking of the status of requests will continue to be useful. Creating transaction records would ensure that all requests are tracked and reported consistently.

The general environment as designed for request of physical materials would apply to electronic document requests, but additional requirements may be identified. To what extent can the Workflow Management Utility as conceived also interoperate with EDD systems such as Ariel?

Electronic scans made in support of ReCAP requests are currently discarded after fulfillment. Workflows could be designed that automatically transfer these scans to HathiTrust. If not already present in the HathiTrust repository, articles and book chapters could be processed through OCR and converted into formats consistent with the repository. While it seems that scans made by ReCAP staff should have similar legal status as those performed by Google, such a step would naturally require careful review. It does seem that the cumulative number of scans performed at ReCAP would represent a significant contribution to the body of digitized library content. This transfer would also need to be accomplished in such a way that does not add additional work to the ReCAP process for fulfilling EDD requests.

Detection of Duplicate Copies of Materials

The ongoing operation of the ReCAP facility will benefit by the ability to identify duplication of materials among the three institutions. Tools can be created to identify duplication among the materials currently in the facility and for new submissions. These tools can be created to be informational, allowing the partners to process duplicates according to prevailing policy or specific circumstances.

The processes designed to extract bibliographic and item information from the three ILS implementations to populate the central discovery index provide an opportunity to identify existing duplicates in the system. This process would complement the de-duplication data and analysis that has been developed through the OCLC collection analysis report and would provide an operational data set that can be used to provide automation support for any de-accessioning that may be carried out. This de-duplication routine would be a subset of the Collection Management module that would be developed as part of the ReCAP support infrastructure. The de-duplication processes will be needed to consolidate bibliographic entries and manage holding data needed for the discovery layer. The data derived from this process can be captured and further leveraged to facilitate collection management, including decision support for de-accessioning.

The programs created to extract records from each ILS and to populate the central index can extended to include algorithms to identify duplicates. Triggers for duplicates might include unique identifiers such as ISBN numbers, but would need many other match points as well. Serials would need to be processed at least at the volume level, if not by individual issues. De-duplication algorithms should be available from other projects. The population of discovery services routinely includes a de-duplication stage making it likely that program code with advanced options and configurability should not have to be created from scratch.

The other and more important consideration involves the identification of duplicate materials as they are accessioned into ReCAP. There are multiple ways in which this process can be handled:

Create a tool to be used in by library personnel as they select materials to be transferred to ReCAP. This tool might include a simple Web interface where the staff member scans in barcode or enters any other unique identifier. Once submitted, the ReCAP Duplication Tool would query the relevant targets:

- Recap Bibliographic database (middleware) (SQL query)
- Each of partner ILS, using Z 38.50 and/or NCIP
- GFA: an additional query, based on the barcode number identified for a likely duplicate could be sent to the LCA system to poll its status. This status would confirm the duplicate or it might reveal an item that has been
- WorldCat (optional, to identify other relevant libraries that hold the item)

After processing the query, the Duplicate Detection tool would return an informational page confirming that the item is unique within the ReCAP partners or listing all the duplicates that exist already in the ReCAP collection or within the partner's active collections.

A batch version of the Duplicate Detection utility would also be created to handle large batches of materials. This material would generate a persistent electronic report, possibly held in a database table, which could be accessed by library staff to make transfer decisions based on duplication status or holdings in partner collections.

The Duplicate Processing utility would be configured according to policies in effect among the ReCAP partners. These polices would control any actions triggered by the detection of duplicates. Messaging options would include notification of both the library considering submitting the item as well as the library that submitted the existing duplicate.

Another alternative might involve extending the GFA software to perform detection of duplicates as new materials are accessioned. Building detection of duplicates into the current accessioning process presents some challenges. First, it might be counterproductive to identify materials as duplicate after they arrive at ReCAP for accessioning rather than making transfer decisions based on duplication status. It is also the case that the inventory data in the GFA software would likely be insufficient to identify duplicates without the support of the middleware to query the ILS systems to retrieve the bibliographic details.

Next Generation Automation

One of the key considerations for the project involves the possible upcoming transitions in the automation environments of the three partner libraries. Based on information gathered to date, it seems that none of the institutions are far enough in planning for next-generation automation systems that short-term development can target anything other than their current environment.

Both of the automation products in the mix, Millennium from Innovative Interfaces and Voyager from Ex Libris have adequate integration capabilities currently to support what is needed for the new ReCAP

technology support environment. Both have been integrated with a variety of discovery environments and have pragmatic means of interacting with discovery and fulfillment services. ILS-DI connectors have been implemented, for example, for both Millennium and Voyager. Both systems support NCIP or have other means to dynamically communicate with services related to patron and item data. Innovative offers optional APIs for Millennium to support patron and Item functionality. The eXtensible Catalog project has developed connectors for both ILS products.

The new-generation systems that the institutions are considering should offer more sophisticated interoperability mechanisms, using native APIs and Web services. Both systems follow a service-oriented architecture that should provide high levels of interoperability.

Both Sierra and Alma, however, stand at a very early stage in their development where the focus of attention has so far focused on basic functionality and not on APIs and interoperability. Princeton, even as a development partner with Alma, for example, has not been asked to test and evaluate APIs that Ex Libris has touted as a distinguishing characteristic of its new system.

In conversations with both Innovative Interfaces and Ex Libris, both organizations report a strategic commitment to delivering APIs that can be used at customer sites to achieve the interoperability and extensibility as envisioned for the ReCAP technology environment. Yet, both companies are currently focused on the basic functionality involved with early production deployments.

It seems clear that short term development will target existing ILS implementations. All work should be done in ways that could easily be adapted to other ILS products and to next-generation library services platforms such as Alma and Sierra.

The same kinds of concerns apply to the online catalog and discovery services in place. Without loading millions of ReCAP records into each ILS implementation, it will be very difficult to support broad discovery in the standard online catalogs of the three institutions. All of the enhanced request features can be accomplished in the online catalogs with discovery limited to the institution's own ReCAP materials. It might be possible to implement a federated search component to expand discovery within these online catalogs, but even this approach is unlikely to be entirely satisfactory.

Further discussion is needed to determine if it is acceptable to offer broad discovery of the ReCAP materials through the discovery interfaces in place (BiblioCommons, Primo, and Blacklight) or if further work needs to be done to attempt to identify mechanisms other than record loading to accomplish comprehensive discovery of the Shared ReCAP collection through the online catalogs.

Future Enhancements / Feature Creep

The proposed architecture of the new generation of technology to support ReCAP should lend itself to additional capabilities. An instance of the discovery service could be implemented that addresses the entire ReCAP collection, even including the non-shareable materials in support of staff functions such as de-duplication and collection development. While the main instance of the ReCAP discovery service would consist primarily of indexing capability to integrate into existing interfaces, this staff-oriented instance would also include implementation of a search interface.

While functionality to support collection development is currently considered out of scope, the new environment is based on an architecture that could be extended to support different scopes. One scope would include all ReCAP materials; another possible scope would span the entire collections of the three institutions. There would be ways to support collaborative collection development, but the design and implementation of the technology platform should scale to the level of a union catalog of the three ReCAP institutions.

Ongoing work

This report serves as an environmental scan and a tentative exploration of functionality and technical implementation options. Vetting of the report among other ReCAP participants may expose omissions or additional functionality, workflows, or technology components that need to be evaluated. A workshop exploring the technology options was conducted on August 21, 2012. Some of the follow-up expected from the partner institutions included feedback on the functionality and general workflows proposed. Prioritization of the modules will also benefit next stages of the project that will develop cost proposals for creation of the technology infrastructure.

Additional research is needed to further refine the technology components that might form the basis of the new environment. Further work, for example, is needed to ensure the feasibility of the index integration strategy proposed, the suitability of Kuali Rice as the low-level service platform, and for Apache SOLR as the indexing platform, and whether the eXtensible Catalog toolkits provide adequate connectivity and metadata management support.

Next stages, performed in collaboration with the Technology Architect, will include the extension of the this tentative design to provide a more detailed expression of the functionality expected, the specific technical components to be used, an inventory of services to be created, and an adequate level of documentation to inform a development team that that would ultimately create the software. This next phase would also include creating estimates of costs of development and implementation and feasible timeframes which would inform the budget of the follow-on project to implement the new ReCAP technology support infrastructure.

Functional and technical requirements need to be developed regarding the work that needs to be done to create a service layer for the GFA Library Archival System. More detailed conversations need to take place with GFA technical personnel. A document needs to be created which more fully describes the work to be contracted to extend the GFA system that can serve as a Request for Proposal that would serve both as a means to determine the cost of the project and to provide the detailed requirements that would drive and assess the development. It is likely that only GFA itself would be in a position to carry out this work given that it is proprietary licensed software.