

Aeon Symposium 2016

“Connecting Users, Pushing Limits”

Opening Keynote Address – Christian Y. Dupont¹

Rackham Amphitheatre, University of Michigan
Thursday, June 16, 2016

*A copy of this talk may be downloaded from my Academia.edu site:
<https://www.academia.edu/26256539>. The text is licensed under a Creative Commons
Attribution-NonCommercial-NoDerivatives 4.0 International License.*



The revolution has started. Aeon users of the world unite!

A spontaneous rallying of Aeon² users had always been my hope, my dream. And here you are! Over eighty of you. From the moment I heard about this symposium³ and was invited to speak, to this moment right now, I’ve been feeling extremely gratified to witness and be a part of this gathering. All the more, since I had *nothing* to do with it. My thanks—and yours—go out to all of the organizers and instigators and mobilizers here at the University of Michigan, who have brought us together in their fun and funky city of Ann Arbor for garden party receptions complete with string trios. Such a treat to be hosted, and treated, by you.

I am also very grateful to our colleagues at Atlas Systems, who have contributed sponsorship support for this event, and moreover have brought us together by bringing us Aeon. I know that the whole Atlas team (half of them are here, it seems) are as excited as you are about our gathering, because the miles you traveled and your presence here show how much you and your institutions care about all the things we can, and want to do, with Aeon, both today and in the future.

Early on, in order for this spontaneous revolution to ignite, we knew we needed to get to a critical mass of users—we figured around 50—like with ILLiad, which had its “nifty fifty” original licensees before Atlas got noticed by OCLC, and entered into an agreement to partner on the development and distribution of ILLiad. And guess what? If you go to the Aeon website and “Look Who’s Using Aeon,”⁴ you’ll see that there are now exactly 50 sites that have implemented, with six more new licensees in the pipeline. So we’ve hit that critical target, and moved beyond it.

As some of you know, Atlas has been hosting an annual ILLiad users’ conference⁵ near its headquarters in Virginia Beach for a number of years now. How many, Dan⁶? 15? The

conference is all about resource sharing, in all its dimensions. I learned so much from those conferences, and each year came away so inspired by the energy and dedication of the hundreds of ILL and access services librarians I met.

In April 2015—just a little more than a year ago—colleagues at Yale University hosted the first user-organized, user-led gathering of Aeon users. It was billed as a “Northeast Aeon Users’ Meeting,” but it attracted a number of you from the Midwest and more or less directly inspired this symposium. That was last year, this is this year. How about next year? Will we be getting together out West or down South for the third annual Aeon Users’ Meeting? Will it become international in scope, like the annual ILLiad conference?

That’s the first bit of anticipatory speculation you’ll hear from me this morning. Lots more on they way. But first, I’ll hope you’ll pardon me if you catch me looking down on at my script. It’s not in keeping with the spirit of revolutionary spontaneity, I know, but I wanted to avoid my tendency become long-winded and digressive—a trait that some of you have experienced and yet have big-heartedly remained my friends. But long-winded you’ll still get, as I do have lots of ideas I’d like to share with you, and so will likely fill my allotted time, leaving time for discussion over lunch or dinner this evening. Still, to keep on track, writing out my ideas seemed the best way to package them for your ears as well as the eventual eyes of others, who may decide to watch the recordings being made of our presentations or to download this text, which I will post on my Academia.edu site this evening.

When I brought the idea for Aeon to Atlas, now more than nine years ago, I regarded it (and still do) as a completely obvious notion. Why wouldn’t we want to take advantage of networked technologies to create a user and request workflow management system for archives and special collections?

From my observations of the world, I’ve learned that the most powerful ideas are those that float around like reactants in a beaker, until some catalyst is added, providing a quicker, alternative reaction pathway to achieve and then move beyond the transition state and rapidly create new compounds. The chemical metaphor is apt an one, and might even bring us closer to what Gandhi might have said when we thought we heard him say: “be the change you wish to see in the world.” For in another context he actually did refer metaphorically to the “invigorating ozone of freedom.” And so we might reformulate the analogy accordingly: “Be a reagent for the possibilities that are worth realizing in the world. Be an enzyme.” There you have it: my little moral maxim for the morning as your stomach works to digest that big bagel you ate on the way in. “Be an enzyme.”

Before speaking about possibilities for the future of special collections user services, which will be my main focus, let me take a few moments to recall the inspiration for Aeon, the problems it

was meant to resolve, and the larger user service and workflow management goals it was designed to support.

Some of you will have heard me tell the story, so I'll beg your indulgence for the benefit of others who have not, especially for those of you here whom I'm meeting for the first time. Again, it is so gratifying to witness a movement that has grown well beyond one person's personal network. For as we'll see, networks in the plural offer the real power and opportunities for growth that spring from the Aeon concept.

I won't go into much detail here, for I've already done so in another talk I was asked to give for California Rare Book School in the fall of 2014, and which I likewise wrote out and put up on my Academia.edu site.⁷

Briefly, back in 2006, so ten years ago, during my first year at the University of Virginia as director of the Albert and Shirley Small Special Collections Library, I became acutely aware of how the semi-automated requesting system that had been developed by our local systems librarians on the back of our integrated library system was frustrating for both patrons and staff. For staff, because it took something like twenty steps to fish out the request and format and print a callslip. For patrons, because in order to request a box of manuscripts, they had to call up the corresponding bib record and then scroll down to the item listing: "Why can't I just request the d—m box from the finding aid!?" There were other limitations, too.

Around that the same time, our colleague Beth Whittaker, then rare book cataloger at The Ohio State University, published an article in *College and Research Libraries* titled "Using Circulation Systems for Special Collections: Tracking Usage, Promoting the Collection, and Addressing the Backlogs"⁸ Beth outlined how she had worked with catalogers, systems librarians, and curators to enable the special collections materials to be checked out in their integrated library system. Alice Schreyer and her teams at the University of Chicago had also worked out a similar process, that likewise included "on-the-fly" item record creation to checkout materials that did not previously have barcodes. Neither approach supported automated requesting from the catalog discovery layer, however, and so were more limited in that respect than our otherwise limited system at UVa.

All this fueled a driving ambition to figure out a better and altogether different approach to the problem, and I would corner anyone who would suffer to listen to me try to think through possibilities. (Kinda like I'm doing with you this morning.) I was feeling pretty stuck until I got my friend Cyril Oberlander⁹ trapped behind a table one day for lunch at Basil's Mediterranean Bistro. Cyril was our Director of Interlibrary Services at UVa. He listened, patiently, and actively, as he always does, and then, when I had run out of breath, told me that the model we should look to for managing special collections users and their requests is not an integrated library system's circulation module, but rather a resource-sharing framework. And with his next

breath he asked if he could introduce me to his friends at Atlas Systems, who happened to be coming up to Charlottesville later in the month to work with him on an extension to ILLiad to support patron-driven acquisitions.

And the rest all happened pretty quickly. Once we had figured out the right paradigm to work with, and the right partner, what needed to be done became patently obvious. With the addition of the proper catalyst, the solution crystallized.

Circulation is about managing the lending of the collection resources of one institution to its own defined user community, whereas resource sharing is about facilitating the discovery and lending of collection resources between institutions. Conceptually, “circulation” happens within a *local* framework, whereas “resource sharing” depends upon and involves the operations of a *network*. The distinction between “local framework” and “network” is an important one to mark, as we’ll be returning to it later as we consider possibilities for devising new types of services around the management and use of special collections and archival materials, and library resources more generally.

The concept of resource sharing involves an expanded notion of user affiliation. In an interlibrary loan request, the user is affiliated with an institution that seeks to borrow an item held by another institution on the user’s behalf. Resource sharing thus depends on a network of institutions that trust one another, and commit to borrowing and lending their resources to one another according to mutually agreed guidelines designed to benefit their respective user communities. It still ultimately involves a closed set of relationships, insofar as institutions and their users either belong to, or are excluded from, participation based on formal affiliations and agreements. Resource sharing nevertheless operates as an inter-institutional network, rather than as a purely internal or local circuit.

One of the chief reasons why a resource-sharing model, rather than a library circulation model, is more applicable to special collection libraries and archives is that our users are not necessarily affiliated with our own institutions. Nor, for that matter, does their use of our collections necessarily depend upon their affiliation with any institutions.

Thus we have a situation that actually goes beyond the expanded but still closed loop of resource-sharing networks into the open environment of unaffiliated, at-large users from the general public. Users at large in an open environment—sounds a bit like a description of the open web, doesn’t it? Once more, we’ll come back to that thought.

An immediate implication for Aeon is the concept that I have dubbed “Aeon Passport.” I was always impressed and secretly thrilled when a participant in a demo I was giving intuitively grasped the notion without my suggesting at all—again, like reactants floating in beaker, ready to fuse. The insight would usually get articulated in the form of a question: “If users go from one

library that has Aeon to another, do they need to create a new account at the second library? Can't they just create one account and use it at every library and archives that has Aeon?" The very formulation of the question reveals that the person asking it is thinking at the network level rather than the local level. Yet there's nothing surprising about that. We're all living out on some tendrils of the web these days.

While we can explain to our users the reasons why Aeon researcher accounts aren't portable, it still isn't very satisfactory to them, nor to ourselves, because we can readily imagine a world in which there would be no such constraint. Yet if we can begin to imagine in it, we can probably begin to build it. When the possible becomes practical, and promises enough value, some enterprising souls will figure out how to make it happen. Again, there are about eighty of you here this morning, so let's pour that catalyst into the beaker, stir it up, and see what precipitates.

What would it take to actually achieve the ideal of portable researcher accounts?

At present, Aeon users must authenticate their login credentials through an authentication system managed by the institution whose resources they wish to consult, and then register their personal information in the institution's Aeon database. Same with their transaction histories, which likewise get logged in the institution's own Aeon database. The database may be hosted locally, or, as has been happening more and more often, remotely with Atlas. Either way, however, each institution's Aeon database remains separate, and there is currently no way for users to view or manage all of their requests at once, or centrally manage their researcher profiles and login credentials.

But what if users could control the authentication process, rather than having the institution do it? Flipping the process around, in other words? We actually do this all the time already, especially on our mobile devices. These days when you download and install a new app on your phone, when the app first opens and prompts you to set up your account, it typically gives you an option to login using your Facebook, Google+, Twitter, or maybe even LinkedIn account.

Now while you may think that allowing a third-party authentication system associated with a social media platform is less secure, think again. It is not only more convenient, but also potentially, or even actually, more secure.¹⁰

When you sign into a site using your social media credentials, the site doesn't keep your password or other information. Instead, you are redirected to your social media login page and authenticate against their servers. The social media credentialing server then tells the app or site you want to access that yes, you are who you say are, and sends you back to the site. This type of authentication is built on a protocol called OpenID.¹¹

Related, but distinct from Open ID is the OAuth¹² protocol, which is used to grant third-party websites and apps permission to access information on another website or social media network.

When used in combination, Open ID and OAuth tell the third-party website or app: you are who you say you are, and here is the information about you that you are willing to share with the new site or app.

With Open ID and OAuth, the authentication and authorization processes are decentralized, and trust is established provisionally with a token system over which the user has ultimate control. It's better than handing over all of your login credentials to every site and every application you interact with, and trusting each of them to keep all your personal information safe.

Early on, Yale worked out a way to allow users to login into Aeon with their Facebook accounts. It was a relatively short-lived experiment because they had to build a local authentication module that piggybacked on Aeon, which became problematic to maintain. Also, back in 2009, the OpenID and OAuth frameworks protocols were fairly young and not fully matured, having been introduced just two years before. Now in 2016, we have available the third generation of OpenID, called OpenID Connect, and the second generation of OAuth. Here at this symposium, I've learned that Yale has recently implemented OpenID to support single sign-on authentication for library applications. There are still some improvements to be made to both protocols, but there is every reason to expect that they will continue to become more robust and even more secure, as there is huge economic interest and impetus to getting it right.

According to a survey conducted in January 2015—so already more than a year ago—four out of five internet users disliked traditional registration forms (gee, I might have guessed a 100%!). Meanwhile, 73 percent preferred to log in to websites and apps with their social media accounts. At the beginning of 2015, Facebook was leading the social media login race with 61 percent market share. It was followed by Google+, Twitter, and LinkedIn, in that order.

Besides decentralized authentication and authorization services, what we really need to make the idea of portable researcher accounts work is a trusted repository of library user data—a repository that also has an application layer that can interface on the one hand with researcher identifier management systems (for example, ORCID¹³) and applications that manage library and archives service transactions (such as Aeon).

Let's walk through how this might work ...

You are a prospective user, and you would like to consult an item held by a library or an archives. You find a metadata record for an item that interests you in an online collections discovery system, such as an online catalogue or a finding aid database. Next to the description of the item, you find a “request” button. You click the button and you are taken to a login screen that asks you to sign into your library request account, or create one if you don't have one. You don't have one, so you follow the steps to create one.

You are then asked whether you would like provide your email address and create a password, or use a social media account to login. You choose the social media option built on the OpenID and OAuth protocols.

You are next presented with a screen that prompts you to consider the extent to which you would like to share your personal information and library-related transactions with other application users and platform services.

Because you have authenticated with your social media account, you may not need to type in very much, if anything. Instead, you may be asked whether to allow your social media account to share certain information it knows about you, such as your name and email address, and maybe your birthdate, not only to be sure that you're old enough to use the library's services, but also for eventual in-person authentication (for instance, so that a staff member can ask you to present your government-issued ID when you arrive in the building as a security check, verifying your name and birthday against the physical ID). Some information such as this may be required by the library's policies in order to provide you with online request services.

At this point, the account creation process could also ask whether you have an ORCID and would like to connect it to your account. Assuming you do, the request application would then interact with ORCID's API and allow you to access your ORCID profile. One of the settings in your profile is a privacy setting that allows you to decide which trusted services you wish to have ORCID share data with, which could include the new portable library user profile platform that we are envisioning.

Now if you don't know what ORCID is, your homework for tonight is to visit orcid.org, and contemplate the ORCID vision of "a world where all who participate in research, scholarship, and innovation are uniquely identified and connected to their contributions across disciplines, borders, and time," and its "principles of privacy and researcher control."¹⁴ And when you're done contemplating, create your own ORCID (it's instant and free) and then allow it to look for and link up any publications you may have via CrossRef Metadata Search, one ORCID's current trusted organization services.

Back now to our own complementary vision for an identifier-enabled research information infrastructure and the process of creating your portable library and archives user profile. To make things simpler going forward, let's give our portable library and archives user profile system a nickname. Let's borrow the name of another Gnostic emanation, like Aeon. Let's call it Nous, meaning "mind." (That's n-o-u-s, spelled like the French first-person plural pronoun, if you're taking notes or tweeting¹⁵.)

At this stage, your Nous account and profile will be essentially complete. You will have securely authenticated with Nous through your preferred social media platform using OpenID and OAuth

services. You will have allowed Nous access to certain information about you from that platform. You will have also linked your ORCID researcher profile to Nous, and given Nous permission to access information about your research interests, affiliations, and publications from your ORCID profile.

The next and last step to be completed in conjunction with your first-time request and account set up will be to give the library or archives from which you are requesting materials or services via its online request management system (Aeon) access to those elements of your Nous profile that the library requires in order to fulfill your request, and allow the library to share request transaction information with your Nous account.

This last step represents a radical departure from the way user and request transaction data is currently managed in an Aeon transaction. As you know, the way Aeon is designed, the principal Aeon database structure includes two principles tables: Users and Transactions. If Nous were a part of the overall request architecture, Aeon would include only a minimally extended User table and a Transaction table that would include link references to your Nous account.

This radical shift would put you in control of your personal library usage data and give you the ability to access and make use of your transaction histories in a variety of ways. You would only have to give up, or rather share control of your usage data to the extent required by the libraries and archives that provide you with services. Libraries and archives would be free to set their own data usage and sharing policies, individually, without respect to other institutions, although one could certainly imagine them establishing a common set of best practice policies and publicizing their compliance with them. You, meanwhile, would be free to set your own policies concerning the privacy of your library usage data, including not only what information you share with libraries, archives, and other agencies, but also with other Nous users and the world at large (more on that last bit later).

As far the library goes, requests from a user with a Nous account would be managed pretty much like they presently are in Aeon, except that certain primary data elements and identifiers would be externalized. As an application, Aeon would still include all of the transactional data and workflow management logic that enables Aeon to do everything that Aeon does.

In fact, Aeon could still manage user and transaction data for users who don't have Nous accounts, and from a practical standpoint, it would of course need to be able to do so until such time as the Nous platform is mature enough, and widely respected and used enough, that it becomes impractical for both libraries and users not to use it—a time when an alternatives to Nous are the exception rather than the norm, and cost users and libraries more time and trouble to avoid them than to take advantage of them.

Think that's a wild idea and impossible goal? Think of the various types of automated transactions systems you already use. Would you really rather carry coins in your pocket to pay a bus or subway fare rather than an electronic pass? Is there a municipality left for which it would be less expensive to maintain a cash-based fare system than an electronic one? In the last year or so, Washington, DC entirely eliminated even paper fare cards on its Metro system. Think about highway tolls. In Boston, as in other cities, there are now toll bridges and tunnels that no longer have toll booths. If you're one of the small percentage of drivers who does not have an electronic plate pass, the toll monitor will take a picture of your license plate and mail you a bill, which you can pay by mail or online. The bill includes a modest surcharge. To cross the Tobin Bridge, for instance, it will cost you \$3.50 instead of \$2.50 if you don't have a plate pass. Once you realize that it's costing you more money and time to pay the bill than setting up a plate pass account, you'll create the plate pass account and be done with it.¹⁶ Same with city parking meters. They're gradually, or, in some localities, very rapidly being replaced by electronic parking meters that take electronic payments. In other places, they're being replaced by pay-by-phone apps. Even more convenient.

Now, I get it. Library use is a fundamental right in democratic societies. So is access to public transportation. And public parking. I'd just like access to library materials to be an easier and richer experience, and one that ultimately costs less to manage and promotes more widespread use.

And I get that privacy of library use is a social value, and one that our ALA and SAA codes of ethics¹⁷ strongly advocate and protect. But frankly, if we're serious about putting our users and their privacy interests first, we should seriously think about putting them in control. Why do we think we know what's best for our users, and offer them only a single, locked-down option in which we own and control all their data through the systems we manage or contract with others to manage on their behalf. Why don't we put users in control, and give them more options and more decisions to make about their own data?

I'll just come out and say it: I think we need to change our attitudes about library patron privacy, including circulation records, which we have tended to hold inviolate and even sacrosanct. We claim to value and therefore protect patron privacy, but what we are really doing is making a choice on behalf of the patron about what is private and what is not. As a patron, I want to own and be in control of my user data. I may want a library to *manage* certain data about me and my use of the library in order to provide me with certain services, but I want to decide whether or not those services are valuable enough for me to divulge certain information. And I also want the library to provide me with an option to easily share my library reading habits with others, just like I do on Goodreads and other social reading sites.

'Nuff said? Hope so. Though I doubt it. I think it will take a little while to win fellow librarians over to this policy position. It may take another user revolution to push us there.

Yes, there will always be exceptions. People who for one reason or another can't use or don't want to use the platform. It doesn't mean that ways of serving them outside the platform can't be offered. But let's not regard the exceptional circumstance as the norm around which we design our services. Rather, let's anticipate what our library users really want, and then build systems that they'll be thrilled and eager to use. So much so that they won't want it any other way.

Tell me, do you meet many users who would prefer to fill our paper registration forms and callslips? Do you have many staff who would prefer to go back to the "old way"? We've already done a good thing by offering them Aeon. Let's keep going, and offer them something even better. Something really great. Something that really helps our users achieve their ultimate aims in using our libraries, which is not just to call up cool stuff to look at it, but to work with the unique materials we are privileged to manage in order to develop new insights into a topic, to generate new ideas, and create new knowledge.

So where were we? Where did we leave off? I created my Nous account in the process of requesting to consult materials from the library—call it Library A. Now I'm continuing to search for information on my topic, and I find that Library B has some relevant material in its special collections as well. So I click the request button and initiate a request with Library B. I'm already logged into to my Nous account, so the next thing that happens is that I get a message on my screen from Library B asking permission to access information about me from my Nous profile and record transaction information in my Nous account. I have reasonable assurance from the way that the request from Library B is presented to me in my Nous account that Library B is a trusted repository. I click "Yes" and continue.

I now get another message on my screen from Library B confirming that my request has been received and will soon be processed. I continue browsing Library B's collections and place a couple more requests, this time with a single button click since Library B and I have already "shaken hands" electronically and agreed to share information with one another. The next day, or maybe a few hours later, or maybe even a few minutes later, I get an email, or a text message, or an account update notification in Nous telling me that the items I requested have been paged and are ready for me to look at in the reading room, or that the digital file I requested is now ready for me view, or download, or pay for—or whatever type of service and mode of fulfillment I requested.

A main advantage of abstracting library user and transaction records from any one system—in our case, Aeon—and placing it on an independent, fully user-controlled library user request management platform—what we have been calling Nous—is that other library and archives software applications can also be designed and configured to tap into it as well, not just Aeon.

For instance, a future version of our main library catalog at Boston College, Exlibris's Alma system, could be configured in such a way as to allow the maintenance of user records in Nous, so that it would no longer have to retain my contact information and details of my circulation transactions. Alma could simply record that a given item has been checked out to a certain user registered with Nous, and when that item is due back. Alma could then periodically check to see whether any library materials are past due, or approaching their due date, or have been requested by another user. In such cases, and only in such cases, Alma would go out to my Nous account and ask it how I preferred to be notified when items are due according to my current settings in Nous. Having previously agreed to give Alma my permission to access this information, Alma would then download it from Nous and use it to prepare and send me an item-due notification.

While this might seem an overly elaborate way to handle a simple circulation transaction, think about inherently more complex transactions, such as an interlibrary loan requests. The way many ILL transactions are handled is still kind of old-fashioned, as in more than 20 years old-fashioned. A user affiliated with Library A finds that Library A does not have a book or article he wants. So he submits a request to Library A to go out and locate a copy and borrow or scan it for him. Library A finds that Library B has a copy and so initiates a request with Library B for it. Library B responds by saying "sure" and then sends the material to Library A. When Library A receives the material it notifies the user that the material is ready. That scenario is known as "mediated" borrowing because Library A plays a primary role in managing the request workflow in both directions.

A more efficient scenario can obtain when Library A and Library B belong to the same consortium and implement an upfront agreement to share materials with their respective users. In this scenario when a user affiliated with Library A wants something that Library B has, the user's request gets routed automatically and directly to Library B, which then automatically fulfills the request and sends the physical item to Library A so that the user can pick it up there, or, in the case of electronic delivery, directly to the user, bypassing the staff and facilities of Library A altogether. This is known as "unmediated" interlibrary loan. It has actually been in use for almost as long as mediated electronic borrowing, but improvements in software systems and communications infrastructure in recent years have begun to dramatically increase the proportion of unmediated requests with respect to mediated requests. The only real limitation at present on the scope of unmediated borrowing is the management of agreements and mechanisms between libraries to provide them.

Unmediated borrowing and lending currently happens between members of closed consortial groups at the institutional level, insofar as institutions constitute the consortiums and support the services. But what if the service were supported by an ILL platform supported in turn by Nous? If a Nous user wanted to place a request for something, he could submit it to a generalized ILL service, giving it permission to read his profile. The ILL service could then check his credentials

and determine which institutions the user would have privileges to borrow from, including, perhaps, repositories that would charge a direct user fee if the user didn't enjoy ILL privileges provided by a sponsoring institution. The generalized ILL service would then scan the available inventories of relevant institutions and place a request with one of them according to an algorithm that takes into account multiple factors when determining where to route the requests—factors such as a delivery costs and speed, and “load balancing” to ensure that participating institutions don't get overwhelmed with requests—the same kinds of algorithms that have already been developed for unmediated borrowing between consortium partners.

From a conceptual standpoint, this model for platform-enabled unmediated borrowing and lending shares many of the same characteristics needed to support the submission and fulfillment of requests for special collections and archival materials, insofar as both scenarios involve the mediation of requests from users who may not be directly affiliated with the institution. And likewise, both scenarios envision and support having users place requests with multiple institutions. Abstracting user and request information from collections information, and managing the two domains separately, at the network level, is the basic, operative idea that would enable the framework not only to function, but to function at a level of scale that would make it viable economically. It would become an essential component of an information management ecosystem that is already large and well-supported financially. It would lift special collections and archives out of their resource-poor ghetto when it comes to systems to manage them. It would put them on the same playing field as other library and information resources, and thereby help to promote their sustainability and use. And isn't that what we ultimately want, for our special collections and archives to be used, and used well, by increasing numbers of people? So shouldn't we invest in means that tend toward that goal, and do so strategically, in a way that makes their management essential and inseparable from other library resources?

Let's talk about other potential benefits of the Nous platform that we're proposing.

Let's say that I have a Nous account and want to share a discovery I made in your archives with other archival researchers. Since my Nous account retains a record of my request, and I have control over that record, I could choose to share information about my request, which is to say, my discovery, with an application that facilitates my sharing such information. It could be as simple as a feature of Nous that packages elements of my transaction into a social media posting. Or it could be an application or service specifically tailored to the academic research community.

Now let's say that in addition to sharing my discovery, I also want to cite it formally. Citations to unique materials really have two components: one bibliographic, the other administrative. As a researcher, I would like a properly formed bibliographic citation that conforms to the style guide of my publication. I also want that citation to include the proper name of the repository that owns the material so that other researchers will know where to go to find it. In the case of citations to

archival materials the administrative information is almost inseparable from the descriptive metadata: a letter from so-and-so to so-and-so, written and dated where-and-when, contained in such-and-such collection, housed at thus-and-thus repository, in box number x, folder number y. In the case of uniquely characterized copies of published materials, say an author's annotated copy, the descriptive part of the citation typically refers to the edition of the publication, which is then followed by the name of the repository, collection name, if applicable, and call number.

What makes citations hard to create and manage is that they have so many diverse elements, and the styles for arranging them differ so greatly. To resolve these difficulties and enable computerized databases to control access to digital publications, publishers worked together to create the Digital Object Identifier standard: DOI.

DOI is, in fact, an international standard, ISO 26324¹⁸ (don't get it confused with that 1970's song by Robert Lamm of Chicago: 25 or 6 to 4!).

DOI is an acronym for "digital object identifier," meaning a "digital identifier of an object"—not, as one might mistakenly suppose, an "identifier of a digital object." The DOI system was initiated by the International DOI Foundation in 1998, and initially developed with the collaboration of some participants in ISO/TC46/SC9, which is the ISO technical committee and subcommittee responsible for developing and maintaining international standards for identification and description, which is to say, some of the core work that libraries and archives are engaged in.

According to the introduction in the official DOI handbook, "the digital object identifier system provides *an infrastructure for persistent unique identification of objects of any type.*" A major purpose and benefits of the DOI system is that it "*enables the construction of automated services and transactions.*"¹⁹

The handbook goes on to explain that "a DOI name is permanently assigned to an object to provide a *resolvable persistent network link to current information about that object, including where the object, or information about it, can be found on the Internet.*"²⁰ The idea is that while information about an object may change over time, its DOI name will remain the same.

We are, of course, familiar with DOIs from article citations. They don't exactly substitute for a bibliographic citation, but as an identifier, they can point you to whatever bibliographic information you need and the content itself. So what else do you need?

Here's a thought: what if were to attach DOIs to all of the items in our special collections and archives? "Crazy talk," like Genie Powell²¹ at Atlas likes to say. There are just too many of them. It would be like Spock's brain. But perhaps we could assign them automatically when processing a request transaction, or attach them retrospectively to our descriptive and item control records. And then, prospectively, we could generate and assign them as we catalog and

process materials, just like we add barcodes to records and container labels. In fact, the only real difference conceptually between a barcode and a DOI is that the barcode identifier is only unique and therefore only useful within a locally contained system, whereas a DOI is unique and useful with respect to the entire internet. Forever.

Remember, a “digital object identifier” is a “*digital* identifier of an object” rather than an “identifier of a *digital* object.” The DOI system provides “*an infrastructure for persistent unique identification of objects of any type*”—think postcard, or pamphlet, or photograph, or folder of photographs. At whatever level we manage our inventory and movement of collection materials for the purpose of describing them and fulfilling user requests to consult or reproduce them, at that same level we can assign a DOI if we so choose, and store it right along with the rest of our descriptive metadata and the administrative information we need to process a request.

Again, quoting from the DOI handbook: “The content of an object associated with a DOI name is described unambiguously by DOI metadata, based on a structured extensible data model that enables the object to be associated with metadata of any desired degree of precision and granularity to support description and services.”²²

So where can we get the all the DOIs we’d need to uniquely and unambiguously identify all the objects in all of our collections? (“All the ganglia, the nerves. There are a million of them. What am I supposed to do? What am I supposed to do? No one can restore a brain!” From Captain’s Log, Stardate 5432.3²³—now there’s a DOI for you!)

Some of you who work with repositories like Figshare²⁴ and Dataverse²⁵ know that you can obtain a DOI for your dataset when you upload it. In fact, those services, and others like them, assign a DOI upon ingest. It’s an integral part of the process.

Now some of you may thinking, why is he recommending the assignment of DOIs to objects in our special collections repositories when he should be talking about ARKs.²⁶ True enough. I just started with what I assumed would be a more familiar concept and coinage, one that has been gaining a good deal of currency. ARKs also are also valid and can be dispensed like cash. In fact, without spending any cash, which is one of their main advantages. DOIs cost somebody something, as there are costs involved in maintaining a DOI registry. Just like you buy ISBNs from Bowker, which is official source of ISBNs in the United States.

ARK (A-R-K) stands for “Archival Resource Key.” Like DOI, ARK is a scheme for the creation and assignment of unique identifiers to support internet services. The ARK schema was initially developed by the National Library of Medicine in 2001 but has since been maintained by the California Digital Library. It has been adopted by more than 450 institutions, including several national libraries, and tens if not hundreds of millions of ARKs have been already been generated. The ARK specification, however, has not yet reached the level of maturity and

standardization as the DOI schema, which is the only such internet identifier scheme that is maintained as an ISO standard.

Also, one potential disadvantage of the ARK scheme for cultural objects, it seems to me, is that the ARK identifier includes as one of its components a numeric code that corresponds to the institution that assigns the ARK. These codes are technically called “Name Assigning Authority Numbers” or NAANs (like the Indian flatbread), which means that the organization has the authority to assign ARKs to objects, typically those that it owns.

Now, what happens when an organization deaccessions an object that has an ARK and that object is accessioned by another organization? It may not happen very often, but it is a realistic scenario that an identifier system has to be able to deal with. Because the ARK is meant to be permanent, the NAAN portion of the ARK will remain the same, along with the rest of the ARK. That’s fine in principle, but in practice a little awkward because someone who looks at the ARK will be led to think that the organization that created it still owns the object.

But can’t a resolver service tell the user that the object is now owned by another institution? Yes, it could. And that points to the need for other components in an identifier system. One of those components is precisely a resolver service: a web-based service that takes an incoming URI and redirects the request to a site that can provide current information associated with the object referenced by the URI (URI: Uniform Resource Indicator—think URL, which is a specific type of URI). A resolver service thus also functions as a registry—a registry of identities and redirects.

The ARK specification describes a “Name Mapping Authority Hostport” (NMAH) and suggests pointing at the Name-to-Thing (N2T) Resolver, which is also now maintained by CDL, which in turn also offers EZID as a service available for an annual subscription that includes creation and maintenance of up to 1 million identifiers per year, for a modest fee.²⁷

Thus, although an organization can register to get its NAAN for free and start generating all the ARKs it wants, to effectively deploy and maintain its ARKs, it needs to subscribe to EZID or another service that supports ARK registration, resolution, and redirection.

In that respect, there isn’t much distinction between ARKs and DOIs. Yet it is important to keep in mind that there are “DOIs” in the plural, meaning the identifiers typically assigned to published works, and “DOI,” in the singular, which refers to the ISO standard and specification.

The DOI standard specifies not only the characteristics of DOI identifiers, but also DOI registration agencies. At present, there are ten DOI registration agencies worldwide, mostly representing agencies concerned with managing scientific publications, although one DOI registration agency, EIDR (E-I-D-R, or “EIDR”), is concerned with unique identifiers for movie and television assets.

According to the EIDR website, “From top level titles, edits, and DVDs, to encodings, clips and mash-ups, EIDR provides global unique identifiers for the entire range of audiovisual object types that are relevant to entertainment commerce. The Registry’s flexible data model supports complex relationships that exist between various assets and is interoperable with other standard and proprietary identifier schemes.”²⁸ Sounds a lot like managing archival assets to me, just within specific content and format domains. The EIDR registry itself has been established as “an industry non-profit that operates on a cost recovery basis.”²⁹ Furthermore, “it is open to all industry participants and offers a variety of membership options.”

EIDR has a governance board, and a technical advisory work group, and graduated fees for “contributor” members. Sounds a bit like ArchivesSpace.³⁰ My reason for mentioning that is not just my own experience as the current member representative for large institutions on the ASpace governance board, but as someone who has observed at close hand how some visionary colleagues in our archives community got together and decided to work out a model for a next-generation archival management system that incorporates the best features of the Archivist’s Toolkit and Archon in a new systems architecture. I likewise believe that some of us could get together and work out in more detail a model for a next-generation approach to managing library and archives user data and request and fulfillment services at web scale—a distributed ecosystem that involves new partners and players in addition to Atlas and Aeon—an ecosystem that includes a user-governed and controlled library and archives user request management platform (what we earlier called Nous) and now also a collections registry service built around the DOI standard (or possibly ARK), which if wanted to continue our Gnostic nicknaming, we might call “Bythos,” meaning “the deep.”

But what would all this get us? What kinds of benefits might the new ecosystem bring us that we don’t already enjoy, or couldn’t readily obtain from improvements to the systems we already have? What should we bother creating new agencies and bringing them into the picture?

Let’s consider a couple of scenarios.

First, assessment. I want to know, globally, what kind of research is being done on a topic, what kinds of resources researchers in the field are finding useful. Applications built on Nous could offer statistical reporting and analysis tools to support a variety of queries on the researcher and transactional data it collects for which Nous users have given their permission to be used for such purposes. It could do so in comprehensive and powerful ways, searching across all users and transactions, or according to selected researcher demographics or other criteria. Nous could then deliver statistical results and even datasets stripped of any personally identifiable information.

Second, researcher connections. As a researcher, I want to connect with researchers who share similar interests, or who might have consulted a specific resource that interests me. An

application could be developed on the Nous platform that would permit Nous users to search for other Nous users whose personal privacy settings permit their profiles and transactions to be searched for matches. If matches are found, the application could facilitate the introduction of the respective Nous users, much the way “friend” or “follow” requests work on social media platforms. The application could provide even more targeted and granular matches of research interests and consulted resources than current platforms like ResearchGate and Academia.edu are currently able to offer, and recommend them to Nous users.

Third—and here’s a really crazy thought—what if we “flipped” collection development, just like we have “flipped” classrooms? The combination of Nous and a DOI compliant collections registry service could enable us to reverse our traditional acquisition and collection development models, whereby repositories assume ownership of archival and special collections materials through purchases and gifts. One of the hardest things about working in a curatorial role, as I now do once again as a librarian, as well as as a personal collector, is realizing just how much interesting and important material is out there on the market or otherwise in private hands: far more than can ever be owned and directly managed by libraries and archives. Researchers, and even fellow librarians and archivists, often don’t realize how limited a view they have of the potential primary informational resources there are to support research. What if we could assist the owners of those materials to record their existence and register their whereabouts—for their personal benefit to be sure, and perhaps the benefit of their insurers, but also to make the materials potentially known to others, and maybe even make arrangements to let them consult them, whether digitally or by arranging a temporary loan—perhaps one mediated by a local special collections library (“flipped” ILL!).

Now if you don’t think that people who own and collect books don’t like to share them, just take a look at LibraryThing. LibraryThing.com presents itself as “a community of more than 2,000,000 book lovers,” whose members have collectively cataloged more than 100 million books. That’s not titles, that’s books. As in personal, individually owned copies. About which they share information, write reviews, make recommendations, and organize into groups—groups of books and groups of readers and collectors who share common tastes.

Founder Tim Spalding, who created LibraryThing as a way of cataloging his own library for academic and bibliophile friends, was a graduate student in Greek and Latin at the time, right here at the University of Michigan. The site launched in 2005, and Tim “had no idea it would explode like it did.”³¹ Too bad, for Tim, as once confided in me that he really wanted to create a cataloging tool for wine enthusiasts and collectors, which probably would have earned him a little more money. But what a great thing he did, and continues to do for the world of books and libraries writ large.

So think about it: what if every book on LibraryThing, which is to say every copy of each book, had its own unique DOI. There would a 100 million uniquely identified, uniquely owned, traceable copies, instantly, and almost none of them owned by any institutional libraries (granted, some smaller institutions have used LibraryThing to publicize their collections and reconstruct historical libraries, but that's more the exception than the norm).

Now what if institutional libraries had a way of uniquely identifying all of the books and manuscripts and archives and artifacts, etc., that they have cataloged? How many billions of items would that represent? And what if they could all be traversed through linked data searches and their provenance traced over time? What if collectors could upload scans of their materials to online repositories that would help them catalog them, and document any known rights information? What if booksellers and manuscripts dealers could likewise adopt a more standardized approach to offering descriptions of their materials, including unique, persistent identifier numbers—DOIs—that would be publicly registered so that buyers and eventual resellers or donors would be able to link additional descriptive and ownership information to the identifier, and decide how much of that information to share publicly. Institutional libraries and archives, could, of course, incorporate these identifiers into their finding aids and catalog records, just like we do with enhanced rare book cataloging records, where we use the MARC 541 field to record references to auction sales and bookseller catalogue purchases. Meanwhile, private collectors may wish to keep their collection records private, just as many of them do now. They would still gain the benefit of improved knowledge and control of their materials, and the ability to resell or donate them the greater ease, since they would be able to readily pass along custodial history and other information in convenient data packages.

I referred a moment ago to linked data. Linked data is the key that would unlock all of these possibilities, and more. To return to our earlier metaphor, linked data can be described as “flipped” data from the perspective of traditional data management models, in which data elements derive their meaning from being part of a defined data set that is bounded and contained—a *database*. Linked data elements, on the other hand, contain within themselves their relations to other data elements and their meaning with respect to entities and collections of entities. What that means is that linked data entities do not need to be contained within a single database, or even related databases, in order to be comprehensively searched. Instead, they can be crawled or traversed along pathways through the semantic web and assembled into a packages of results, such as knowledge graphs—those handy profiles of places and people that appear in the right-hand column of your Google search results.

Think of the knowledge graph you could create for a rare book or manuscript if various linked data entities could be traversed and threaded together on the basis of their sharing a common digital object identifier, or DOI. Such knowledge would be dynamic and cumulative, accruing

new layers of meaning each time someone interacts with the object and leaves a trace of that interaction through the addition of some new linked element or entity somewhere on the web.

What I'm after here is a next-generation approach to managing library and archives user data and request services—a distributed information architecture that depends on more players than just your institutions and Atlas and Aeon. Atlas can build or build out some of the pieces with Aeon, but other critical parts don't exist yet and must be built by other agencies that, in my view, should operate as not-for-profit, membership-governed, and user-driven organizations.

It is important to recognize that the architecture I'm proposing would not be specific to special collections and archives. It could support other types of libraries and library services, and even independent collectors, booksellers, and other agents, as I have suggested. This is a crucial point to grasp, because in order to build something of the magnitude that we're proposing, significant resources would need to be engaged at the outset, and a sustainability plan developed around a market forecast that has a realistic chance of meeting the plan's targets. Sustainable business models have been developed for other services upon which library and archives communities have come to depend, plans that involve both not-for-profit and for-profit partners, such as the one that sustains ArchivesSpace, which includes fees collected from Registered Service Providers, like Atlas, as well as member institutions that use the software.

Business models and plans can, and I hope will be developed to catalyze the transformations of our services we have envisioned here this morning. To bring those to a crystalline point: the idea is a simple, and I hope obvious, one: we need to massively externalize user data and collections data in order to build a library and archives that is truly of the web, by the web, and for the web. A library and archives available to everyone, everywhere, all the time.

The revolution has started. The reactants are swirling in the beaker. Let's all be enzymes, and precipitate the change.

NOTES

All urls were accessed on or around the date this talk was given.

¹ Formerly Aeon Program Director for Atlas Systems; since September 2014 Burns Librarian and Associate University Librarian for Special Collections at Boston College; see: <https://bc.academia.edu/ChristianDupont>

² See: <http://www.atlas-sys.com/aeon/>

³ See: <http://www.aeonum2016.org/>

⁴ See: <http://www.atlas-sys.com/aeon/who-is-using-aeon/>

⁵ See: <http://www.atlas-sys.com/conference/>

⁶ Dan Specht, CFO/VP for Sales and Marketing, Atlas Systems: <http://www.atlas-sys.com/about/address-staff-directory/>

⁷ Christian Y. Dupont, “‘Take Things Always by Their Smooth Handle’ And Give Them a Good Turn. Reflections on Life and Libraries,” talk presented for California Rare Book School at The Book Club of California, San Francisco, November 3, 2014: <https://www.academia.edu/9188322>

⁸ Beth M. Whittaker, “Using Circulation Systems for Special Collections: Tracking Usage, Promoting the Collection, and Address the Backlogs,” *College & Research Libraries* 69/1 (January 2008), 28-35. doi: 10.5860/crl.69.1.28.

⁹ Cyril Oberlander is currently Dean for the Library, Humboldt State University; see: <https://www.linkedin.com/in/cyril-oberlander-a7a5987>

¹⁰ See, for example, Fahmida Y. Rashid, “Signing Into Websites With Google, Facebook is Good for Security,” *PC Magazine* Digital edition, May 21, 2015; see: <http://www.pcmag.com/article2/0,2817,2484486,00.asp>

¹¹ See: <http://openid.net/>

¹² See: <http://oauth.net/>

¹³ See: <http://orcid.org/>

¹⁴ See: <http://orcid.org/about/what-is-orcid/mission>

¹⁵ For tweets from the 2016 Aeon Symposium that used the official #AeonUM2016, see: <https://twitter.com/hashtag/AeonUM2016?src=hash>

¹⁶ See Martine Power, “Drivers cheer cash-free tolling on Tobin Bridge,” *Boston Globe* June 21, 2014; see: <https://www.bostonglobe.com/metro/2014/07/21/cash-free-tolling-tobin-leads-faster-commute/waF280CLp1KZjwLzyCEaxJ/story.html>

¹⁷ For the Code of Ethics of the American Library Association (ALA), see: <http://www.ala.org/advocacy/proethics/codeofethics/codeethics>; for the SAA (Society of American Archivists) Core Values Statement and Code of Ethics <http://www2.archivists.org/statements/saa-core-values-statement-and-code-of-ethics>

¹⁸ See: http://www.iso.org/iso/catalogue_detail?csnumber=43506

¹⁹ DOI Handbook, Introduction: https://www.doi.org/doi_handbook/1_Introduction.html#1.5

²⁰ DOI Handbook, Introduction: https://www.doi.org/doi_handbook/1_Introduction.html#1.5

²¹ Genie Powell, Chief Customer Officer, Atlas Systems: <http://www.atlas-sys.com/about/address-staff-directory/>

²² DOI Handbook, Introduction: https://www.doi.org/doi_handbook/1_Introduction.html#1.5

²³ From the Star Trek episode “Spock’s Brain,” original airdate September 20, 1968; see: <http://www.chakoteya.net/StarTrek/61.htm>

²⁴ See: <https://figshare.com/>

²⁵ See: <http://dataverse.org/>

²⁶ For information about ARKs (Archival Resource Key Identifiers) referenced here and below, see: <https://wiki.ucop.edu/display/Curation/ARK>

²⁷ See: <http://ezid.cdlib.org/home/pricing>

²⁸ See: <http://eidr.org/>

²⁹ See: <http://eidr.org/about-us/>

³⁰ ArchivesSpace is an open source archives information management application; see: <http://archivesspace.org/>

³¹ See LibraryThing.com, “Who we are”: <https://www.librarything.com/howweare.php>