

Background and Goals

After expenses from a frugal first workshop, the funding that remained from the original award was used to partially support a follow-up workshop, called Registering, Discovering, and Using Distributed Services in Academia, as an Advanced CAMP. (See <https://spaces.internet2.edu/display/servint/Advanced+CAMP>.) One of the outputs from the original workflow workshop was that scholarly and scientific workflow need to be placed in the overall context of loosely coupled scholarly and scientific service architecture, and that key issues in workflow, such as registration and discovery of services, warranted additional, broader examination. This workshop was a direct consequence.

The follow-up workshop was held on June 19-20 in Minneapolis, attended by a large and deeply engaged group of participants. Key Mellon-funded projects, including Bamboo, SEASR, Fedora, and Quali all sent lead architects who both presented and participated in extensive discussions. Other important open source elements of the application/middleware space in R&E, including Internet2 projects (Shibboleth, Grouper, etc.), grids (OGSA) and policy/security expertise all engaged. Valuable corporate participation included IBM and Microsoft.

The focus of the workshop was on the basic issues – service registration, discovery and use – presented in the challenging use cases of scholarly services. Typically, scholarly services are loosely coupled, either hosted in facilities on campus outside central IT or even more distributed to partners in other institutions, countries, or market sectors. The format of the workshop consisted of presentations by projects (SEASR, Bamboo, Fedora, Quali, OGSA, etc) on their use of service architectures, followed by discussions by service infrastructure providers (ESB deployers, IBM, and Microsoft). Extensive general discussions then wrestled with the forest amidst the trees and break-out groups addressed

Data models – Is formal data modeling important enough that taxonomies and controlled vocabularies be put in place before coding begins? What development methods lend themselves to adaptable data models? Who needs to play in this space, both for intracampus services and for interrealm services?

Data governance – Who owns the data in a service-oriented approach? Can the role of data steward evolve to deal with external data needs and users? Are federations the vehicles for working out interrealm data sharing agreements? Who should be at the table?

Registering and Discovery – What means are used to facilitate service discovery on campus? Can they be extended to work with external, loosely coupled services? How are service gaps or overlaps identified and dealt with? Who needs to be at the table?

Findings and Outputs:

It was an intense workshop, with a daunting task of structuring discussions and analysis in such a new and complex area. Several themes cut across all the findings below, and many use cases contain issues on both sides of the scale:

- Compile-time versus run-time issues – For example, determining if the license of a service permits its inclusion in a package is done far in advance of identifying a running instance of the service in execution
- Local versus hosted – For example, significant issues in privacy, HIPAA and FERPA arise when a service instance is located outside an enterprise, while there are significant issues in economics in operating some services, such as scientific computing, locally.
- Machine versus human consumption of meta-data – Service description and service discovery requirements of people and computing processes are quite different, particularly in the arts and humanities.

1. There is an extremely wide set of meanings to “scientific and scholarly services”, which in turn greatly affects the basic issues of registration and discovery.

At the workshop, the spectrum of what was defined in use cases as a service was quite broad. The primary dimensions were in scale and complexity. For some a service was a well-known API to a specific and atomic function; for others a service was a rich and unique scholarly capability, read about in a journal or conference. The breadth of meanings was reflected in discussions around human-readable versus machine-readable discovery (and the need for both in many uses cases), and in automatic versus manual and vetted registration. Service description is another significant, highly textured construct also affected by scale. For some projects, service description is parsable XML and for others it is a one-page document with some standard boxes for developers to enter text into.

Yet another distinction was in whether service meant the abstract or an instance. Often there is need to discover if there is a service in the world, and then determine which instance to be used. In library terms, it is the distinction between an item record and a bibliographic record. Sometimes, for robustness, policy or other driver, there is a need to do a traversal from an instance to the more general abstraction in order to find another instance to use.

Clearly there would be good value in building common terminologies and categorizations that would help structure discussions and developments over the next several years.

2. The transition of SOA from enterprise service bus to loosely-coupled services is profound.

All aspects of SOA are transformed by the move to loosely coupled scholarly services. Regardless of whether the “looseness” is intra-campus inter-unit, interinstitutional, or international, most existing SOA issues are expanded, and new ones such as regulatory compliance of services emerge. For example, storage might normally just be a service, but when that service is loosely coupled, by law Canadian universities must know the physical location of that service before invoking it.

From discussions on distributed services governance, it was clear that registration implied some sort of certification or reputation system. There was occasional remembrance of Gopher not just as an Internet navigation system but as an implied vetting process for the “validity” of the sites in the tree. (There was also occasional remembrance of several Tom Lehrer songs.) Some discussion pursued the merits, for scaling, of domain or discipline specific registration for loosely coupled services in a “Dewey Decimal” hierarchy. However the issues of discovery of cross-hierarchy services, especially for infrastructure services such as IdM, was one of several challenges identified to such an approach.

Many of the impacts in the shift from intrarealm to interrealm services are in its policy consequences, and are covered in item 6 below. In interrealm use cases, service discovery is quite an art. There are a variety of techniques, and an application developer or enterprise might have to deal with several. DNS SRV work but are rarely used. XRI and XRDS enable resolution services but each has significant problems as well.

As privacy takes on greater definition, the notion of attaching consent (specifying what permissions the user has given for any information or attributes that they release) is becoming central, requiring that data and services that consume the data need to understand the consents. Such trans-realm requirements are present in many scholarly use cases. An OASIS standards group (XPSA) is now being formed to address the interoperability of consent in at least the health care sector. Since service API's today do not deal with consent, transitions will be tricky.

While the issues in loosely coupled services are as much as about registering and discovering services as using them, the traditional ESB concerns of transformation, orchestration and transport are still needed and current ESB technologies might be extended to be used interrealm once loosely coupled registries and discovery are better understood.

There is a critical relationship to note between shared data and common practices that permit the effective and secure use of the data, especially in inter-domain instances. Lacking corporate governance or regulatory definitions, nurturing common practices is essential if scholarly and academic services are to share data across realms. Whose role this is – professional associations, academic federations, or some other – is quite unclear, but the lack of common practices will be the greatest impediment to shared data in the next few years.

Maintaining the authenticity of data in a loosely coupled environment is an unknown art at this point. Provenance issues abound. Moreover, given the nature of sequencing services, does each service that does a transformation create a new scholarly object? Must this be signed and registered and its provenance recorded.

3. Computing as a service

How to provision computing, and to a lesser degree, data, as a service was a continuing thread throughout the workshop. Two approaches dominated the discussions – clouds and grids. Within grids, there are several meaningful subsets, such as high-performance, high-throughput, etc. A number of papers and presentations (see both the Gannon citation and the IBM talk on the wiki) sharpened the distinctions between the approaches, the costs and benefits of each model, the sets of applications that work best in each environment, and, increasingly, the legal and policy implications of the different approaches. Today, the architecture, or worse the specific instance of the architecture, determines the way computing is presented as a service, and there is little externalization by the computing service provider of functions such as IdM to other services, nor is there interoperability or portability among flavors of Grids and clouds. Data as a service was declared even beyond our advanced scope right now.

4. Applications and developers

There is still uncertainty around how application development, particularly in the scholarly areas, taps into both enterprise cyberinfrastructure and loosely coupled services. Numerous projects mentioned the gap between the increasing complexity of the infrastructure and the domain application developers, particularly in the social sciences and humanities.

One interesting discussion noted that the patterns of invocation of services and access to data are often very useful to others. Much as in on-line shopping, where “customers who looked at X also looked at Y” referrals can be helpful, there could be value in some cataloguing of how services and data are searched for and used. At the same time, there are sensitive issues about privacy, IPR, and chains of credit that need to be carefully handled.

5. Technical Issues

The growing diversity of the service technologies platforms, combined with the expectation that loosely-coupled scholarly services will use multiple approaches, creates a daunting set of technical issues. The variety of technologies at play in SOA - services vs. web services coupled with the variety of approaches within each (REST, etc...) make integration of loosely coupled scholarly services quite difficult. Several technologies at several different layers must be profiled when attempting to assemble an integrated scholarly offering. It was noted that publish and subscribe is a paradigm, not a single standard. The variety of specific implementations here also deeply affects interoperability. Because of the overall complexity in scholarly services, the value of an enterprise meta-architecture greatly increases.

Registries are critical, but a single universal one won't service the community. There will be multiple single registration points. Particularly in the arts and humanities, service discovery is a human oriented, often serendipitous process; though linking the service at a code level ultimately needs to happen. Metadata, persistent and stable as services move, is critical in these processes.

Rather than opening up legacy commercial applications to open institutional infrastructure, there was a sense that many large application vendors use a SOA as the new point of lock-in. Coupling external or loosely coupled services to embedded vendor SOA seemed difficult. Though we did not have representatives from the large open source enterprise service bus systems such as Mule, there was the sense that using external services might take significant reworking of those packages as well.

There are performance issues about provisioning computing and large scale data as web services. Passing XML across HTTP is just not efficient; its only virtue is that it goes through firewalls. There is also a marked distinction between high performance and high throughput computing, and how they are provisioned differently as services.

There is a growing trend for using the enterprise IT organization not only as a provider of services, but as an evaluator of external services for local users, and integration of those services with campus infrastructure, even though central IT may not operate the final product.

6. Policy

Numerous issues in scholarly services suggest the need for a more direct engagement with enterprise IT. The value proposition for enterprises to use services outside their control (either loosely coupled on campus or in external domains) needs to be articulately made, and address enterprise privacy and security concerns. This is particularly true of the HIPAA and FERPA consequences of privacy spills. “Reasonability” is an important concept for managing data and service requests within the enterprise; it needs to be extended for scholarly and loosely coupled services. There also needs to be some audit mechanisms about external applications using the service. This is both for documenting value of a service in an academic setting and for managing the data with proper administrative controls. Lastly, there are a set of issues, and remedies, around external services transiting enterprise firewalls that need to be better documented. The answers lie as much with the service developer as with the enterprise firewall deployer, but both must be engaged.

Many of the issues in policy created by scholarly SOA are “compile-time”. For example, there are often concerns about whether the license associated with a service is consistent with the other licensed services being used by an application. Such considerations tend to occur in human-oriented service discovery processes and are best addressed at compile-time. Small but significant differences in licensing have complicated the use of external services. E.g., Quali uses ECL as its license, but is making adaptations to Apache as part of the Quali effort and must use the Apache license for those adaptations.

It was noted that there was growth in both identity federations and in federated repositories. Alignments between those two efforts would help user experience and

perhaps reduce legal overhead in building inter-institutional agreements. There is some natural tendency in that direction but more might be done.

SLA's around services is an important but completely unexplored area. The first requirements for SLA's are emerging in the commercial cloud provider space, but at some point scholarly services, particularly computing, will need to consider SLA's.

There was extensive discussion of emerging privacy activities and how they might affect academic SOA. EU privacy directives, coming out of the Article 29 WG, now require much more specific behavior. Moreover, the EU has now stipulated that service providers outside of the EU must obey EU laws if operating within the EU. This affects US campuses and content services as service providers and US campuses as identity providers. Scholarly services in the US will soon need to be aware of the issues.

Next Steps and Fundable opportunities:

Several next steps could provide considerable benefit:

- A fuller characterization of the issues in loosely coupled scholarly services and some promotion of common terminology and categorizations would be quite helpful. This characterization could also identify areas for further investment. A related resource would be some best practices in building and using services. There was some discussion of the role OKI should or could play.
- The technology of scholarly services is far out in front of policy. There is obvious value in some effort to motivate campus and enterprise interest in scholarly services. A set of activities should encourage policy discussions about the many issues identified above that arise from distributed scholarly services. Distributed academic services will require articulation of the culture changes to key constituencies – CIO's, developers, lawyers, etc. It seems reasonable to leverage the attention that Bamboo and Kualu are getting to discuss these policy issues
- There should be a concerted effort to push services infrastructure across major open source applications. For example, a linkage between enterprise services and Kualu library services should be established early. Several clusters of services are particularly important, such as identity management services, which are needed by almost every vertical application and done well by few of them. (The attendance of Duke central IT at the CAMP, and their considerable understanding of enterprise and federated infrastructure seems serendipitous for Kualu Library.)
- More focused discussions should be initiated to maximize the alignment and overlap of identity federations and federated repositories. This responsibility might be assumed by the group of international R&E federations.

- There is merit in considering a domain-centric approach to service registration and discovery. Can the taxonomies of scholarly or discipline communities (ACM, MLA, IEEE) be useful in organizing services, including controlled access? Although it is quite early in our practice, it might be useful to host a discussion that brought together scholarly service developers with curators and others in taxonomies and folksonomies to inform issues in service description, registration, and discovery.
- Use of standard profiles for services would reduce the deployment variations and increase integration of services with each other. These profiles are technical specifications of community best practices – a summary of how campuses assemble the service infrastructure for service based scholarly applications. Such profiles could significantly increase the adoption by new campuses in scholarly services participation.