A Brief History

This document is a brief history of the development of Internet identity within the R&E community. After a short discussion of the "Internet identity" term, it attempts to describe the activities within the R&E community and their influence on other sectors.

By Internet identity, we mean the set of mechanisms that allow the use of local authentication and attributes outside the perimeter of the organization that provides the local authentication and attribute services. "Authenticate locally, act globally" is a shibboleth of this vision.

Background:

It is important to note that the same factors that led to the pivotal role of R&E in Internet development also motivated the R&E community to take a leadership role in the creation of Internet identity. Those factors include:

The need for researchers to collaborate between institutions
Distinctive privacy requirements that shape well-engineered approaches
Close interactions with the government sector and agencies that drive R&D investments
Strong relationships with standards organizations and corporate research groups
An orientation to open-source solutions

During much of the 1990's, identity was contained within each application. Authentication and most attributes were done by the app. However, by the late 90's, the idea of a central authentication service for the enterprise had taken root. The major one was the development of Kerberos by MIT. Its creation for Project Athena and subsequent adoption by Microsoft created a foundation that organizational identity and Internet identity could be built on top of. Another major precursor was PKI. As a technology, PKI offered global trust, at the expense of easy local deployments for both technical and policy challenges. A third major motivator was the plethora of applications emerging as the web developed and organizations could run multiple local applications to provide services to their users. This need for single sign-on, implemented via a web browser, drove the creation of systems such as CAS, a popular open-source software system, and PubCookie.

At the same time, activities hosted by Educause engaged the R&E community with the federal government on a variety of PKI efforts. There was a continuing theme of establishing an R&E PKI infrastructure via a sector certificate authority. The approach would be similar to what the federal government was building with their PIV cards, agency-based certificates, and using the federal PKI bridge that was already providing trust mappings between federal agency infrastructures to potentially bridge to the R&E PKI root. The discussions never led to significant deployments, but they provided valuable lessons about where the flexibility and pivot points in trust needed to be for Internet scale identity. It was the soil for federated identity to grow in.

The Beginnings:

Much of the leadership for the R&E community work in identity came from RL "Bob" Morgan, an architect first at Stanford and then at the University of Washington. Through his vision, his interlocking personal relationships with members of key standards organizations such as IETF and Kantara, and his ability to articulate complex thoughts with simple mumbles, he was pivotal in shaping much of the work until his untimely passing in 2012, and even thereafter. Capturing a list of other key individuals would only do injustice to those many accidentally overlooked in the process.

There were several organizational vehicles for much of the work. Chief among them was a set of middleware activities within Internet2. The organization was focused on advanced applications and therefore constructed networks to support the applications. It became apparent that while the networks and the applications were advancing, there was a layer of middleware missing that was necessary to enable the authentication and authorization for those applications. That need led to the Internet2 Middleware initiative that was a developmental hub for the work in Internet identity. Much of the work was done in close partnership with Educause, a broad higher ed IT community, with their expertise in outreach and training. Federal agencies, in particular NSF and NIH, provided vital grant activity that supported the work, and those same agencies became adopters of the work product. Another important group was "Stone Soup", a set of leading universities that supplied much of the intellectual capital and early proving grounds for middleware work. A set of ad hoc activities connected these US efforts to expertise in Europe, where R&E networks in Sweden, the UK, France and Spain, among others, had the same drivers as in the US for federation. That locus eventually expanded to include Japan and Australia.

Internet identity was also being worked in some pockets of industry. IETF had long been the home for PKI standards, and bar BoF's there had explored many of the issues in building Internet scale identity. OASIS became the standards organization for SAML. Companies such as IBM and MS contributed expertise to the Internet2 middleware efforts.

The First Work: Directories and Schema

The initial efforts in the R&E community, beginning in 1999, were around directories and schema. It was evident that directories were critical to a scalable middleware infrastructure, and that a shared schema among the R&E community would enable meaningful sharing of attributes. Important principles such as differentiating values in the directories from values on the wire, and using privacy-preserving attributes and identifiers, helped set the foundations for federated identity. Reference documents such as the LDAP Recipe and the eduPerson schema were developed, along with community consensus processes, heavily influenced by IETF, to create community standards.

The placemat of federated identity:

The major driver for much of the Internet2 Middleware work was to create a sustainable and scalable community mechanism for privacy-preserving inter-realm authentication and authorization. Use cases in science, in scholarly activities, and in community service drove a set of requirements that were the basis for investigations into using PKI, extending the Kerberos PAC concept to inter-realm use, and other architectures. After those evaluations, it was decided that each approach had significant problems, ranging from deployability to scalability. The decision was made to pursue strategies that were extensions of institutional websso systems, pushing the concept of single sign on into a general inter-institutional infrastructure.

That decision led to a number of key consequences. It added a unique set of privacy consequences. It charted a multi-lateral approach, and with that, critical use of metadata. It led to trust frameworks that would be both locally deployable and globally scalable.

In early 2001, an architecture for interinstitutional web sso was established. Echoing an early IETF design principle that a good protocol should fit on a cocktail napkin, the architecture was sketched one night on a restaurant placemat. Because of the breadth of use cases it had to address, the approach was more complicated than first envisioned, and so there was a six month pause for a reduction effort, attempting to simplify the flows and reduce the number of components. The architecture that emerged from that review, around the end of 2001, was the original, but the insights and understandings developed in the review led to a more rapid and solid development effort.

Around the same time, OASIS was beginning to convene a WG around a security markup language to address use cases of outsourced applications working with enterprise identity infrastructure. Key R&E identity leaders met with the OASIS business community to determine how to partition development activities. The business use cases were all bilateral relationships, while the R&E model was multilateral. It was agreed that the basic bilateral exchange protocol – SAML – would be developed within OASIS, with that architecture designed to accommodate multilateral requirements, such as supporting scoped attributes. The multilateral dimensions – from scoped attributes to shared metadata and schema – would be addressed within the R&E community development. That distinction and the resulting collaboration (many of the writers and editors of the OASIS SAML draft were those also developing the R&E multilateral aspects) was perhaps the most critical factor in the success that followed.

From the R&E development came several enduring elements of Internet identity infrastructure. The first was Shibboleth, a widely used and definitive open source federation software system, in use across a variety of sectors in the US and internationally. The second was a sequence of federations such as InQueue, that resulted in InCommon, pioneering trust models that are both deployable and scalable. The federation work in turn drove yet another major component of Internet identity, end-entity metadata.

Sequencing the building of institutional trust was delicate. The lessons of failed PKI deployments indicated that rigid policy requirements needed to be replaced with flexible and “roughly consistent” community agreements. Pivot points were placed in key locations to allow local adaptations as long as they were documented and available. Progressively over several years, the community joined InQueue and then moved to InCommon as they were able to document their trust-related procedures. This scale through flexibility was intended to be replaced over time by scale through consistency as the value of multilateral federation increased via the network effect. And progressively, that consistency has grown and baseline expectations are being established.

The rise of R&E federations worldwide:

Other countries had the same R&E drivers for interinstitutional identity and, starting in the early 2000’s, a few places in Europe, most notably the UK, the Netherlands, Switzerland and Scandanavia began to experiment with federated identity. The relatively small size and high expertise allowed rapid development and simplified architectures. In particular, hub-and-spoke models became effective national level approaches, with the central hub taking on a variety of roles including being the federation operator, acting as an attribute release and consent point and being the IdP for its members.

While the growth was in national level federations, the R&E community is global, and inter-federation was always the desired end-state. In Slaughter England, a first international gathering was held (informally called “Leading Trust to Slaughter”) in 2004 to both develop strategies to encourage federations in other countries and to build enough consistency in approaches to facilitate eventual inter-federation. One output of that meeting was the eventual formation of REfeds, an international organization of R&E federations.

Other R&E Internet identity activities:

The federation work was not the only activity harnessing the R&E community. There was a steady stream of ongoing work in PKI. This included establishing a Higher Ed Bridge Certificate Authority (HEBCA) with innovative technology approaches at Dartmouth and a policy authority with broad R&E representation. Another activity was to establish a US Higher Education Root Authority (USHER) that would be include in campus versions of web browsers to provide inter-institutional capabilities. In addition, for ten years, NIST, NIH and Internet2 hosted a yearly workshop in Gaithersburg on cutting edge PKI and other trust issues. The sessions, and the hallways around them, were quite valuable in facilitating academic and business cross-pollination.

Attributes:

Though the early emphasis in Internet identity was the strength of authentication and LOA (levels of assurance), the R&E community, with its privacy and access control use cases, had a particular focus on attributes and identifiers. This was evident in starting the eduPerson attribute schema well before starting software development to exchange those attributes. The consensus process took quite some time, balancing a rich set of requests with the understanding that deployments became harder the broader the schema and interoperability became less likely with wide attribute controlled vocabulary.

There were some particularly useful engagements in this area. A "Tao of Attributes" workshop at NIH in early 2007 was one of the first investigations of the relevant metadata about attributes that were both viable to generate and responsive to needs. One outcome of that workshop as a cautionary tale about the swamps of LOA of attributes. Early conversations were also defining about the scoping of attributes and effective ways to extend values within a community of interest. Building semantic communities of practice became useful as InCommon grew as a schema proving ground.

Closing comments:

While the Internet identity work done in R&E can be considered a true success with profound consequence on the current state of Internet identity, a case can be made that a full set of components needed for the R&E use cases have not happened. In this sense, the work is an immature success. There is certainly ongoing some maturation in some aspects, such as the progression towards consistent trust approaches via baseline expectations. But the overall environment has been limited by the quick and widespread adoption of the basic approaches while key pieces of the original vision remain unimplemented. It has become its own embedded base. Examples include the lack of good consistent tools for user to manage their privacy and attribute release, the inability for relying parties to signal required and optional attributes in a fine-grain fashion, and the slow emergence of good IdP discovery approaches.

Withal, there is a global, functioning internet identity infrastructure, a layer of connectivity that enhances the basic Internet network layer.