Interfederation and Metadata Exchange: Concepts and Methods

Ian A. Young
SDSS
ian@iay.org.uk

Chad La Joie
SWITCH
chad.lajoie@switch.ch

V1.10, 20 May 2009

1. Introduction

In the last few years, a number of identity federations have been established around the world. Although most of the federations familiar to readers will be those established in the domains of higher and further education, the idea of federated identity is now beginning to take hold more widely. This is resulting in both the growth of established federations beyond their original communities as well as the formation of many federations serving entirely new communities.

This paper is written in the context of federations which use the SAML protocols and metadata specifications. Although the technical work that will be described here will initially be specific to this environment, we believe that the conceptual framework is applicable to all federation technologies.

To date, individual federations have been developed, and grown, largely in isolation. This has led to great, and desirable, diversity within their technical and social structures: indeed, it is probably true to say that each extant federation has a different definition of the word “federation”, one that is in line with the goals of the particular community served by its creation. In this paper, we enumerate some of the different components of the various definitions of “federation” so that we may discuss concepts related to federation organisations with more precision.

As federations have grown, the communities they serve have inevitably begun to overlap with each other, leading to increased interest in the idea of “interfederation”. As with “federation”, this is a word where different component ideas can be combined to form many different definitions. In practice, most people assume a definition of “interfederation” that is strongly grounded in their own definition of “federation”; for example, someone whose definition of “federation” includes strong behavioural guarantees between members is likely to see this as an essential part of “interfederation” as well.

We believe that the diversity seen in existing federations argues against using these existing models as a general approach to interfederation. Instead, we propose a clear division between those aspects of individual federations that are specific to a particular community’s focus and those aspects whose technical nature allows more standardisation across federations. Our working definition of “interfederation” is, as a result, based on the concept of an enabling technical infrastructure for federated identity at internet scale, in the same way that existing federations provide such an enabling technical infrastructure as *one* of the services they deliver today to their own community. We anticipate that this new infrastructure will initially be operated by existing federation operators, but that in the long

term some of these operators may choose to “move up the stack” in order to focus on the needs of their particular communities once the required technical infrastructure has become available as a commodity.

Finally, we describe in outline form methods by which such a global technical infrastructure may be built around the exchange of entity metadata between actors that we call metadata registrars, aggregators and publishers. Federations and other organisations can participate in this system of metadata exchange in many ways depending on the combination of roles that they choose to adopt.

As part of the outline, we describe a number of specific technical developments on which such an infrastructure would depend. In some cases the technical community has already started to develop components of this infrastructure.

2. Federations and their Roles

We define a federation as a member organisation serving the federated identity needs of some particular community. This definition is deliberately broad in order to allow for the observation that across the many such federations that exist today, there is a wide variation in the the selection of services each provides to its members. For example:

- Some federations establish legal obligations between their members, others rely on members to reach bilateral arrangements where appropriate.
- Some federations act as purchasing consortia, others exist solely to enable communications between members.
- Some federations require deployment of particular software in particular configurations, others merely provide recommendations.
- Some federations collect and publish metadata on behalf of their members, others depend on third parties (in some cases, other federations) for such technical infrastructure.

These, and other, variations are significant, widespread, and inevitable given that federations exist to serve different communities which have different underlying needs. This situation has little in the way of practical consequences for a single, small federation; however it becomes critically important when federations grow and begin to border on each other – or even overlap – and interoperation between members of different federations is desired.

An early response to this situation, during a period in which a wide variety of identity protocols was in use, was to devise inter-federation strategies in which federations would deploy protocol gateways to convert from one protocol to another and incidentally provide connectivity between their members. The reduced variety of protocols in use today makes the protocol conversion aspect of these proposals less relevant to most federations. In addition, a solution involving high-layer gateways ignore the Internet’s central architectural “end-to-end principle” and as a consequence suffers from a number of unsatisfactory security properties.

Most current federations deal with this area by regarding it as an implicit requirement that two communicating parties belong to the same federation, and require prospective users

to find a federation which is prepared to accept both parties as members. One result of this has been that some organisations (to date, principally service providers rather than identity providers) have been required to register the same entity with many federations.

One way of describing this state of affairs is that federations are currently thought of as containing their members, the entities belonging to the members and the communications between those entities, with the implication that membership of a particular federation is a prerequisite for communication. It is important to note, however, that such a requirement is not a technical one: federated identity software such as Shibboleth is not functionally aware of the social and legal constructs we call "federations". The only technical requirement for two entities to be able to communicate securely is that they be able to acquire trustworthy metadata about each other. A single federation common to both entities is one way to achieve this – and an excellent help to scaling within a small community – but far from the only possibility.

The requirement that entities have access to trustworthy metadata for other entities with which they communicate allows them to establish *technical trust* in each others' identities. Technical trust is insufficient to allow entities to interoperate in other than a very restricted environment, however, as knowing the identity of a potential communications partner does not establish its good intent. Such *behavioural trust* can only be established through real-world agreements between the owners of the entities involved. Behavioural trust is entirely dependent upon technical trust having been established, but is otherwise separate; we can therefore view behavioural trust as being *layered* on top of technical trust.

This separation and layering of trust concerns is a characteristic of other Internet-scale technologies. For example, purchasing from an on-line bookstore involves the technical trust acquired through the TLS protocols implemented by the user's browser and the store's server. This allows the user to establish that the site being viewed is the one intended and not an impostor. It does not, however, establish behavioural trust in the store's good intentions: for this, the user must rely on a real-world relationship of some kind, perhaps established through an explicit contract or by laws related to the sale of goods.

Both technical trust and behavioural trust can always be established directly between two parties: for example, by the direct exchange of metadata and contracts respectively. For a very small number of partners, this is often the simplest way to proceed. At larger scales, however, the number of potential exchanges grows as the square of the number of participants. Federations exist primarily to ease this consequence of scaling by acting as a *broker* of trust of both kinds on behalf of the members of their community:

- Behavioural trust is brokered by the establishment of behavioural norms appropriate to the community represented by the federation. This may be entirely informal, in the form of recommended practices to which members may assert compliance, or a formal part of a legal agreement members make either with the federation operator or multilaterally across all members.

- Technical trust is brokered by orchestrating a process through which trustworthy metadata is made available to federation members for the entities with which they wish to communicate. This may in practice involve any combination of registering, aggregating and publishing metadata on behalf of members: the essential component for the members is that the federation provides access to a source of trustworthy metadata

for their use on which they may then layer behavioural trust. The principal value added by the federation here is to allow individual metadata consumers to delegate the process of determining which sources of metadata are sufficiently trustworthy.

Early federations attempted – to varying degrees, and with varying degrees of success – to conflate these two concepts. We believe that the success of federated identity at Internet scale requires, on the contrary, that these layers be clearly separated so that they can both evolve away from the "federation as container" model in appropriate, and different, directions:

- Behavioural trust is community based and will therefore be likely to remain associated with federations as member organisations. As existing federations grow, however, their communities will necessarily lose coherence. We expect the result to be a model in which the set of participants are partitioned, not along national or regional lines, but instead in smaller, more community-focused federations which may have more complex relationships, often being nested within or otherwise overlapping with other federations.

- Today's geographically-based federations provide much of their value by brokering technical trust between their members. It is possible to extend this model through formal bilateral arrangements to exchange metadata in a limited way between federations. However, such an approach is likely to have only short-term success, both because formal bilateral agreements scale badly and because, as indicated above, the hermetic federations are unlikely to survive in their present form. We believe instead that a more natural long term evolutionary path leads instead to a "metadata layer" model. This would have properties analogous to other Internet-scale services such as the Domain Name Service, e-mail and the World Wide Web in which participants' expectations are that information registered anywhere is freely available to all interested parties.

Although this separation of concerns leads to the possibility that individual federations might choose to participate to different extents in the brokerage of trust of different kinds, we expect that for the foreseeable future most federations will continue both to establish behavioural norms and to register and publish metadata.

3. The Metadata Layer

We propose an Internet-scale "metadata layer" for federated identity with characteristics similar to those provided by such services as the Domain Name Service as described above.

Entity metadata originates with the real-world *owner* of the entity, who interacts as a *metadata registrant* with a *metadata registrar* in order to register that metadata for use within the metadata layer. The registrant's expectation will be one of *universal publication* of the metadata as registered and this should be allowed for in any agreement between the registrant and registrar. Anything less than potentially universal publication of metadata would constrain the usefulness of the particular registrar to an extent that registrants would prefer to register elsewhere.

Each metadata registrar operates under a set of procedures which we suggest should be documented and published as a *registration practice statement.* The procedures under which a particular metadata registrar operates will vary depending on the effort the registrar is willing to undertake in order to verify and maintain each registrant entry. We believe, however, that it is possible to divide the contents of a registrant's metadata into two kinds of element: those whose values are critical to the security of relying parties, and all others which do not fall into that category. The number of such security-critical values is surprisingly small and easy to enumerate, and we expect that most metadata registrars will be able to agree on appropriate practice for their verification.

Metadata registrars will almost always also be *metadata publishers,* providing access to the metadata they have registered to *metadata consumers* using some *metadata publishing protocol.* The simplest form of this is a well-known location from which the aggregated metadata can be retrieved, with a digital signature included to provide an integrity check and to ensure that the retrieved metadata can be associated with the publisher and thus with the appropriate practise statement. The signature has no other implicit semantics; in particular, the presence of a signature does not imply that the signer makes any general representations as to the truth of statements made in the metadata, other than as specified in the practise statement.

The situation where a metadata registrar also acts as a publisher for its own (and only its own) registered metadata, and where members are both metadata registrants to and metadata consumers of that metadata, corresponds to most existing federation operations. The extension of this into a metadata *layer* requires the introduction of another actor, which we will refer to as a *metadata aggregator,* and whose role is to aggregate metadata from multiple sources.

This role is implicit even in the single-federation case already discussed: federations today act as registrars, aggregators and publishers of their local metadata collection. In the metadata layer, however, the emphasis is on aggregation across sources, each representing a different registrar or indeed other metadata aggregators.

A metadata aggregator should, like a metadata registrar, operate on the basis of a set of published policies which we will refer to as an *aggregation practise statement.* This statement is required in order for consumers of the aggregated metadata (whether end entity owners or other aggregators) to make a judgement as to whether the aggregator's choices of sources to aggregate meet their needs as relying parties. It is in short the mechanism by which trust in a particular aggregator's output can be made *transitive* across the metadata layer to the original metadata registrar.

In summary, in a metadata layer model, a metadata registrant registers metadata with a metadata registrar according to the registrar's practice. The registrar publishes that metadata for the benefit of, amongst others, aggregators who act as trust intermediaries by combining metadata from multiple registrars whose practice statements they find acceptable. After passing through some (probably quite short) chain of such aggregators, metadata is published for the benefit of ultimate metadata consumers by some federation's designated metadata publisher, whose own practise statement gives those consumers assurance as to the processes by which metadata has been accepted as trustworthy by their federation's standards.

In the early stages of development of a metadata layer, we believe that it is likely to grow through bilateral relationships between existing federations. However, several use cases have been proposed in which groups of federations acting as metadata registrars may come together with a third party, which does not operate a separate registration function, acting as a regional aggregator across the federation grouping.

4. Metadata Layer Technology

The technology to operate a metadata layer at Internet scale does not yet exist.

Although metadata registration is a function of most existing federations, it is performed in very different ways in each case and we do not expect this situation to change. Instead of trying to standardise registration, we observe instead that existing federations publish metadata in very similar ways (as simple signed aggregates hosted at well-known locations) and propose devoting effort to the development of a proof-of-concept *aggregation engine* capable of performing the aggregation and re-publication functions required by a metadata aggregator, initially of SAML metadata in particular.

Although in general we believe that it is essential that the metadata layer be viewed in the long term as a “common carrier” of metadata, we observe that significant variation exists in local conventions and requirements among federations today. The pilot aggregation engine implementation will therefore need to be able to perform arbitrary filtering and transformation operations on metadata being processed.

When existing federations act as metadata publishers today, it is usually through a well-known location from which a single document containing all the metadata can be downloaded. Such a scheme is adequate for small numbers of entities, but clearly inappropriate for the publication of metadata for very large numbers of entities as would ultimately be available through a universal metadata layer.

We therefore propose the development of additional metadata publishing protocols to provide access instead to individual entity metadata documents. In particular, we intend to develop a specification for a RESTian “pull” protocol that allows end entities and aggregators to deal with metadata documents for individual entities while preserving the model that these are part of an enclosing aggregate.

In the longer term, we believe that a move from “pull” protocols to a mixture of “pull” and “push” will be required for a number of reasons, such as reducing the latency between a metadata change being made by an entity’s owner and the availability of that change to consumers. The additional implementation and deployment complexity required is significant, however, and we believe that implementing “push” protocols is unnecessary for initial proof-of-concept development.