ABFAB
Internet-Draft
Intended status: Informational
Expires: September 12, 2013

J. Howlett
R. Smith
Janet
M. Wasserman
Painless Security
March 11, 2013

Trust Requirements in a Federated World
draft-howlett-abfab-trust-router-ps-03

Abstract

TODO: This document outlines the requirements for trust in a federated environment, and enumerates the requirements for a trust infrastructure. It also examines existing trust infrastructures given these requirements and concludes that none fulfil all of the requirements, and suggests that maybe a new one is required that does.

Table of Contents

1. Introduction

Trust is a concept whose exact definition differs depending on the exact scenario in which it is being applied, however, there is a large degree of commonality about the meaning across all contexts - the idea of "trust" usually centres around two entities (be they people, organisations, or whatever) establishing confidence and faith in the reliability, truth, or abilities of each other.

In the world of federated technologies, trust typically means exactly that: two federated entities being able to establish confidence in each other. What this means more specifically is that two entities are able to verify that each other is who they think they are (i.e., that it is not another entity pretending to be someone else), that they represent the organisation they say they do (i.e. an entity isn't misrepresenting themselves), and that transactions between the two can be secure, unaltered, and verifiable.

This document examines the requirements of a trust infrastructure in this context of the federated world. It then looks at existing trust infrastructures, examining their pros and cons in relation to these defined requirements. It then draws some conclusions about work needed to bridge some missing gaps.

2. Terminology

This document uses existing terminology that the reader may not be familiar with - and also introduces some new terminology - so a small set of definitions is now included.

- Authentication Policy Community (APC): A set of entities that share a common trust infrastructure and a common set of identification requirements for the entities to be admitted to the community.

- Community of Interest (CoI): A set of federation-capable entities who want to interact with each other for a common purpose and with a specific set of requirements.

- Entity: A general term for IdPs and RPs.

- Federation: An agreement between various entities that allows for delegation of trust based a common set of semantics between an RP and an IdP.

- Identity Provider (IdP): An entity that asserts identity information about its principals to RPs.

- Principal (or Client): An individual (i.e. a real world person) or a computer that is registered with an IdP and is attempting to get service from an RP.

- Relying Party (RP): (a.k.a. Service Provider (SP)) The federated entity representing an organisation that consumes identity information about a principal asserted by an IdP in order to provide a service to that principal.

- Trust Arbitrator: A central point of trust infrastructure that gathers and passes on crowd-sourced trust recommendations about members of its APC. The entities within the APC provide trust recommendations to the arbitrator: the arbitrator does not make trust recommendations on its won. The entities in the APC trust that the Trust Arbitrator will ensure that recommendations from the APC membership are correctly reflected in the trust rating presented.

- Trust Advisor: A central point of trust infrastructure who entities rely on it to (make trust recommendations about other entities and?) trust decisions about who is admitted to a CoR based on a set of advertised criteria.

- Web of Trust: A mechanism for extending trust between two different trust infrastructures where an agreement exists that each of the trust infrastructures can use the judgment of the other infrastructure to make its own trust decisions. A Web of Trust can be transitively extended across multiple trust infrastructure. In other words, you and all your friends become my friends.

3. Trust

3.1. What is Trust?

Trust is a word and a concept that can mean many things depending on the context in which it is being discussed, and can encompass a wide range of requirements. For example:

- In personal relationships, trust between two friends is usually a mutually established relationship where each can rely on the other for confidentiality and their help in times of need.

- In airline travel, trust between a consumer and an airline provider is largely a one way relationship where the trust of the airline by the consumer means the consumer expects the airline to keep their aircraft well maintained, to get them to their destination alive and (roughly) on time, and to (try) not to lose

their luggage.

- In e-commerce, trust between a consumer and a vendor is also largely a one way relationship where the trust of a vendor by a consumer means that the consumer considers the vendor to be likely to provide good service, a reasonable price, and to fulfil their order after it has been paid for.

Many such examples of trust exist in all walks of life and across many contexts. Across most - if not all - of these, some commonalities appear in the meaning of trust. Trust usually centres around two entities being able to prove to each other who they are, and then use that as a basis to transact in some form in a manner that is confidential and secure. Indeed, the OED defines trust as "Confidence in or reliance on some quality or attribute of a person or thing, or the truth of a statement".

3.2. Communities of Trust

Trust, as defined above, is not something that just appears from nowhere - it must be established somehow. The simplest way is, of course, for two entities to build it bilaterally through a process of slowly increasing the level of trust in each other through a series of transactions over time. This process can be expedited, however, through one of three commonly employed methods:

1. Trust can be established transitively through commonly trusted 3rd parties who have previously built up a level of trust with each other (i.e. a web of trust, such as is seen in the PGP web of trust)

2. Trust can be bootstrapped indirectly by a Trust Arbitrator - an entity that is trusted by all who make use of it to appropriately manage community provided reviews of trust (e.g., recommender systems such as eBay member ratings)

3. Trust can be bootstrapped directly by a Trust Advisor - an entity that is trusted by all who make use of it to appropriately manage directly established trust relationships between it and its community members (e.g., X509 Certificate Authorities).

In practice, each of these mechanisms can be, and are, used in different contexts. They are typically deployed on a per-community basis, such as the eBay community, a research and education SAML federation, PGP users, web site deployers, etc. - because in a world of billions of people and many millions of organisations, each of whom is a member of many communities, finding a single entity, or a small number of entities, that are commonly trusted by all to be a

Trust Arbitrator or a Trust Advisor is simply impossible. So, as a way of tackling this scaling issue, "Communities of Trust" have organically appeared in all three guises of expedited trust establishment: sometimes in a decentralised way (e.g. PGP key management through a web of trust), sometimes arbitrated by a Trust Arbitrator (e.g. Trust management between sellers and buyers on eBay), and sometimes by a Trust Advisor (e.g. X509 CAs, Twitter "verified accounts", SAML federation metadata managed by the federation operators, etc).

Within the world of computing, Trust Advisors are a typical way of establishing trust, as it is essentially a method for people or organisations to outsource the problem of trust establishment. However, to achieve this, someone or something needs to set up and manage the Trust Advisor, and users of it typically need to pay for the service. Consequently, this method of solving the trust establishment problem is largely used in business to business and business to consumer communities. However, for those communities where no financial transactions are involved, this method of trust establishment is less appropriate as the community members do not always want to - or are unable to - pay for the service.

Trust Arbitrators are less commonly seen, and are usually found where a Trust Arbitrator stands to make financial gain through the use of its services - either directly (e.g. members pay for its trust recommendations) or indirectly (e.g. a chargeable service having more members because of the increased trust these members can have in each other). Again, this method may not be appropriate for those communities where financial transactions are not involved.

The Web of Trust method of establishing trust is least seen of the three methods, since while it is decentralised and appropriate for any type of community, including those where no financial transactions are taking place - transitive trust can be built up in a way without paying for it - it is the hardest and most time consuming way of establishing trust.

3.3. Authentication Policy Communities vs Communities of Interest

In those scenarios where Trust Arbitrators or Trust Advisors are the methods of bootstrapping trust, two separate communities actually exist. However, the two are often conflated as they are typically equal, by accident of design. These two communities are:

- The Authentication Policy Communities: the set of entities who are known to a trust infrastructure and for which the trust infrastructure has a notion of trustworthiness.

- The Community of Interest: the set of entities who want to interact with each for some particular purpose (e.g. to do B2C to B2B transactions, to collaborate as part of a cross-organisational project, to communicate securely, etc).

These two types of communities are often conflated because of the lack of Trust Arbitrators or Trust Advisors that span multiple communities of interest. This results in each community of interest using a single (or small amount of) Trust Arbitrators or Trust Advisors, thus the Authentication Policy Communities and Community of Interest are one and the same.

In the real world, however, there are many cases where communities of interest want to span multiple Authentication Policy Communities. For example, SAML entities in the education sector who are part of different Authentication Policy Communities (i.e. SAML Federations) that are often geographically split (each country typically runs its own SAML federation) often desire to communicate; [TODO other examples].

When entities from different communities of interest want to transact in this way, trust has to be established across communities. This can achieved in one of three ways:

1. Entities can join multiple communities, if the technology allows this.

2. A Trust Bridge can be set up between Trust Arbitrators/Advisors to allow entities from different Authentication Policy Communities to communicate.

3. A Trust Arbitrators/Advisors can attempt to become the arbiter of trust for multiple communities.

The first of these options is a rather inelegant way of solving the problem, and will likely involve significant organisational effort per community the entity joins, and is therefore not a particularly scalable solution, and is therefore only a workable workaround in limited circumstances.

The second option scales much better but requires significant effort by the Trust Arbitrators/Advisors in ensuring that their rules of registration are compatible.

The third option has significant drawbacks, however - either the Trust Arbitrators/Advisors has to relax its rules so much to be inclusive for a range of communities of interest that it becomes very limited in the assurances it can offer, or it has to impose a set of

standards that many entities will not be willing to opt into due to the high costs it imposes on them to meet these standards.

3.4. Trust in a federated environment

Given the general view of trust presented above, what does trust look like more specifically in the federated world?

In a federated world, there are a few types of entities and trust relationships, to whom trust each means slightly different things. There is:

- Principal to IdP, IdP to principal. Trust here is mostly organisational around an existing relationship between principal and their IdP. The principal trusts the IdP to control its data properly and to assert correct information about them. The IdP trusts the principal to keep their credentials confidential so that the IdP can authenticate the principal when required.

- Principal to RPs, RPs to principal. Trust here mostly flows through guarantees between the principals and its IdP, and that IdP and the RP, that allow the user to gain access to services offered by the RP.

- IdP to RP. Trust here centres on the IdP and RP being able to verify each other's identities, and often associate that identification with an existing business relationship between the organizations.

So in a federated world, there are actually two types of trust in play. Firstly, there is technical trust (i.e., is the server I'm talking to really the one I think it is, and are all communications secured?); and secondly there is organisational or behavioural trust (i.e., the trust that allows two entities to know for sure that that the other entity represents a particular organisation that they have a business relationship with, what guarantees are in place with each other about their relationship, etc).

The second of these is a real world problem to be addressed and not a technological issue and thus needs to be solved out of band. It might, of course, have technical components to it (e.g. schema definitions so the two understand the semantics of the data transferred back and forth), but those would be part of the out of band negotiation.

The first, however, is at the core of federation. The next section looks in more detail at what federation needs from a trust framework to properly establish technical trust.

4. Trust requirements in a federated world

The requirements of technical trust in a federated world largely centre on entities being able to verify each others identity and establish secure communications channels to allow for signing, encryption, etc. To achieve this, various properties of the trust infrastructure are required. Some of these requirements are absolutely critical, while some are optional requirements to help with scale. These are detailed next.

4.1. General Requirements

First of all, we will look at the main general requirements.

4.1.1. Identifying Partners

In most cases, it is necessary for organisations to have confidence that the configuration that they have for their partners actually corresponds to their partners and is not, for example, an attacker claiming to be their partner. This is largely an issue of vetting the claimed identity of organisations when they join a community.

An additional requirement here is that organisations need to minimise the cost of validating their partners' identities, and of proving their own identity to their partners.

4.1.2. Connecting to Partners

This is probably the most obvious requirement. Organisations must be able to establish trust with each other through configuration, i.e., servers representing a particular organisation must be configured to be able to connect to servers representing partners of that organisation in a secure verifiable manner.

To be able to connect, the configuration must include such things as specifying the protocols to be used and the cryptographic operations to be used.

Additionally, to be able to effectively connect to partners, organisations must have effective tools for managing policies that control the flow of information between the two partners.

4.1.3. Clearly Delineate Registration from Usage

As detailed earlier, the conflation of Authentication Policy Communities vs Comunities of Interest hinders progress of large scale adoption of federation and transitive trust establishment. Conceptually splitting those two out gives many advantages, not the

least of which is the possibility of being able to better support communities of interest that span multiple Authentication Policy Communities. Authentication Policy Communities will provide the technical trust, while Communities of Interest overlain upon one or more Authentication Policy Communities can provide the behavioural trust.

4.2. Specific Requirements

Alongside the general requirements presented above, there are some more specific requirements. These are discussed next.

4.2.1. Many IdPs, Many RPs

Entities must be able to establish trust with an arbitrary number of partners - scale should not be an issue.

4.2.2. Frequent changes in Community Membership

It must be possible to support changes in membership (adding new partners, removing former partners, or changing the configuration of partners) with minimal operational effort, and without requiring manual configuration changes that could result in new partners having delayed or incomplete access to services, or former partners retaining some access to services beyond the end of their membership.

Additionally, organisations want to minimise the cost of managing these frequent changes - there should be no requirement for credential exchange between a large number of parties, no manual configuration changes on a large number of systems, etc.

This is a requirement for both Authentication Policy Communities and Communities of Interest, though the scale of changes is likely to be different in the two.

4.2.3. Costs incurred by community that benefits

Typically in federation, the operational cost related to adding and removing partners is born by the RPs, who need to maintain bilateral credentials for each IdP whose users can access the services provided by the RP. This is fine in a case where a single RP provides services to a group of IdPs that pay for membership in the community, or pay for access to specific services.

However, in a less-centralised partnership the costs of exchanging credentials with each IdP could serve as a disincentive for organisations to provide services to the community and/or result in cases where an RP is unwilling or unable to incur the costs of

providing access to new partners. Therefore, if a trust framework supported a mechanism where the operational costs are distributed to the organisations that are receiving benefit from incurring the costs, it would help drive adoption and usage.

4.2.4. Minimal Costs/Effort for forming new communities of Interest

To expand the use of federated services, it should be possible for a group of potential partners to form a new Community of Interest with minimal infrastructure and the lowest possible operational expense.

In order to minimise start-up costs, it should be possible to leverage existing shared credentials and use those credentials for a new Community of Interest.

Practically, this resolves to two problems:

- It must be possible to create a new Community of Interest that uses credentials from one or more existing Authentication Policy Communities.

- It must be possible for a partner to join multiple Communities of Interest using a shared Authentication Policy Community, and for different entities (such as users or servers) within a partner to participate in different Communities of Interest. Practically, this means that information about the Community of Interest in use needs to be transmitted to an IdP, so it can be used as part the authentication process.

4.2.5. Flexible Communities

It should be possible for Communities of Interest to grow to encompass more partners, partners in different regions of the world, or partners who have different Authentication Policy Communities available to them.

It must, therefore, be possible for a single Community of Interest to be serviced by multiple Authentication Policy Communities. While it might be necessary for any given RP/IdP pair to share at least one Authentication Policy Community, it should not be necessary for all of the partners within a given Community of Interest to share a single Authentication Policy Community.

4.2.6. Flexible Trust Links

Related to the above, it should be possible for Communities of Interest that span multiple Authentication Policy Communities to be able to support that transitive establishment of trust over these no

matter what technology is used for trust establishment in each Authentication Policy Communities. This will help increase the flexibility of communities.

4.2.7. Multi-Role participation

It must be possible for a single partner to participate as both an RP and an IdP within a single community (either a Community of Interest or a Authentication Policy Communities).

4.2.8. Multi-Purpose communities

It must be possible for a particular Community of Interest to be equivalent in membership and configuration as an Authentication Policy Community. A use case for this requirement would be an Authentication Policy Community that provides services to its own customers, perhaps for maintenance of their own Authentication Policy Community membership.

5. Analysis of existing Trust infrastructures

TODO: Intro to this section. N.B. Much more detail needed, should be a decent amount of analysis. Not a huge amount though - severa paras per tech, perhaps.

5.1. PKIX

TODO: A trust advisor (or small set of trust advisors) in the form of heirarchically signed X509 certs. Only Authentication Policy Communities, not CoI. Costs incurred by the service, not the user. Works well but only when governance and management is done properly. Which it isn't, generally. Can devolve trust establishment a bit with intermediate certs.

TODO: Not particularly suitable to federation in the real world as existing X509 heirarchies are usually not run well enough, and running one well enough is a significant cost.

5.2. PGP style web of trust

TODO: Web of trust style trust establishment, obviously. Works well but not massively deployed because of the technology understanding required of users to sign each others keys and whatnot.

5.3. SAML Metadata

TODO: Can use PKIX, or directly established public/private keys through metadata, to establish trust. In the former case, all of

PKIX drawbacks. In the latter, needs a trust advisor controlling the SAML metadata. Usually conflates APC with CoI (though entity categories allow for CoI in a very limited sense). Doesn't scale well because of it! Interfederation as a point solution to scaling having to be designed.

5.4. FreeRADIUS shared secrets

TODO:

5.5. Other Things?

TODO:

6. The Future of Trust

TODO: Trust is something that needs engineering - trust frameworks and systems are needed that fulfil all of the requirements stated if we want a way of establishing trust between two artibrary federation entities that cross arbitrarily large and diverse communities, and in a way that empowers the actual users of the systems to make best use of it all. No particular solution fulfils all of the requirements above that are necessary for wide adoption of federation in a way that is easy and cost effective for the communities using it. Any new solution should consider all of these needs.

7. Acknowledgements

TODO: The document authors would like to thanks Jim Schaad and Klaas Wierenga for providing review.

8. Security Considerations

TODO.

9. Privacy Considerations

TODO.

10. IANA Considerations

This document does not require actions by IANA.

11. References

11.1. Normative References

[OASIS.saml-profiles-2.0-os] Hughes, J., Cantor, S., Hodges, J., Hirsch, F., Mishra, P., Philpott, R., and E. Maler, "Profiles for the OASIS Security Assertion Markup Language (SAML) V2.0", OASIS Standard OASIS.saml-profiles-2.0-os, March 2005.

11.2. Informative References

Authors' Addresses

Josh Howlett
Janet
Lumen House
Library Avenue
Harwell Oxford
Didcot OX11 0SG
United Kingdom

EMail: josh.howlett@ja.net

Dr. Rhys Smith
Janet
Lumen House
Library Avenue
Harwell Oxford
Didcot OX11 0SG
United Kingdom

EMail: rhys.smith@ja.net

Margaret Wasserman
Painless Security
356 Abbott Street
North Andover, MA 01845
USA

Phone: ++1 781 405 7464
EMail: mrw@painless-security.com