Network Working Group
Internet-Draft
Expires: August 29, 2013

M. Shore
No Mountain Software
K. O'Donoghue
ISOC
February 25, 2013

A problem statement on trust in IETF protocols
draft-shore-trust-problemstatement-01

Abstract

This document attempts to set out a problem statement and framework for future discussions regarding "trust" in the IETF.

Table of Contents

1. Introduction

"Trust" is used quite broadly in IETF documents but has not been discussed or defined very rigorously. To the extent that it's been discussed explicitly it's typically been within an implementation or protocol definition context, often around the question of trust anchors and their management (see RFCs [RFC5914], [RFC5934], [RFC6024], and many others for examples).

In this document we intend to tease out how IETF protocols have tended to approach questions around trust, discuss whether or not this has been sufficient, and see if there is new work on trust that could be of value. We are not specifically interested in defining the word "trust," but rather identifying broader issues and problems related to trust.

Note, as well, that a survey of trust mechanisms in IETF documents and protocols is out-of-scope for this document.

Text relating to problems around revocation will be added to future revisions of this document, as well as text relating to problems modeling trust in third-party and federated authentication and authorization protocols.

2. Terminology and glossary

Assurance, Assurance Level

Attestation of Control

Authentication

Binding, Cryptographic Binding

Blocklist/Whitelist

Certification, Certification Practices, Certification Practices Statement

Certifying Authority, Certification Authority

Confidence

Correct

Digitally Signed/Digital Signature

Hijack

Identity, Identity information

Leap of Faith

Legitimate

Mediated Trust

Revocation

Risk

Source Integrity

Transitive Operations (in the context of Trust)

Trust

Trust Anchor

Trust Auditing

Trust Establishment and Bootstrapping

Trust Framework

Trust Revocation

Trust Passing

Trust Transaction

Trustee, Trustor

Unilateral Trust, Bilateral Trust

Validation, Validation of Compliance

3. What is trust?

As of this writing, "trust" does not appear to be defined in an IETF document, relying, rather, on functional or operational contexts to imply intent. [RFC4949], a quite substantial internet security glossary, does not define it anywhere, although in its definition of "source integrity" it includes the text:

> The property that data is trustworthy (i.e., worthy of reliance or trust), based on the trustworthiness of its sources and the trustworthiness of any procedures used for handling data in the system.

which is a rather circular discussion.

Kaliya Hamlin, on her IdentityWoman blog, has written a bit [1] about this in the context of the NSTIC [2] (National Strategy for Trusted Identities in Cyberspace) program, ultimately arguing that a "trust framework" is an accountability framework.

Accountability would seem to be a component of an intuitive understanding of "trust," and establishing accountability is a core component of establishing trust as we understand it, but accountability implies auditability only; that is to say, this definition seems to focus on the ability to retrospectively determine that some set of actions took place in the past. Establishing accountability does not establish that future interactions will be safe, and authorized.

The OASIS Web Service Secure Exchange Technical committee has developed a standard [ws-trust] to specify a framework for, among other things, brokering trust relationships. Much like the IETF, they seem to rely on an operational definition of "trust":

> Trust is the characteristic that one entity is willing to rely upon a second entity to execute a set of actions and/or to make set of assertions about a set of subjects and/or scopes.

Zainab Aljazzaf has done extensive work on trust in web transactions, and in his doctoral dissertation [tbss] he defines trust as a complex subjective term, with the following components:

- Utility: a trustee (for example, an IdP, or a web server) needs to provide a utility to a trustor (for example, a relying party or a web browser)

- Dependency and reliability

- Risk attitude.
- Vulnerability
- Remedies, in the event of a breach of trust. This is closely related to Hamlin's notion of accountability
- Confidence expectation, with an inverse relationship between trust and confidence (Aljazzaf asserts that possession of confidence makes trust unnecessary)
- Context-specific
- Subjective -- trust is experienced differently for different trustees
- A trustor may have no control over the trustee. The more control a trustor has over a trustee, the less need there is for trust

He then goes on to propose the following one-sentence definition of trust:

> Trust is the willingness of the trustor to rely on a trustee to do what is promised in a given context, irrespective of the ability to monitor or control the trustee, and even though negative consequences may occur.

The key question seems to be around risk, and the expectations of risk. Imparting trust would seem in some sense to be a signaling mechanism - that transactions with the trusted party should be regarded as carrying known than transactions with non-trusted parties.

4. Modeling trust

Describing, or modeling, trust requires identifying parties and understanding their relationships. It may also require identifying trust processes.

Participants in a trust transaction may resemble the participants in identity services (see, for example, the OASIS SAML 2.0 glossary [samlgloss]). A trustor may be seen to take on a similar role to that of a relying party, while a trustee may have a parallel role to an identity provider. Both a trustee and an identity provider make authoritative assertions about a subject.

Trustors may or may not have an established relationship with a given entity. That is to say, at the time that a transaction is initiated, a trustor may have information about the other party, and they may have an existing business relationship. For example, if you have an account with your bank, you provide them with identifying and other information. When you establish an account with an ISP you are providing them with considerably less information but you are paying them - you are purchasing a service.

It is also common to have unilateral relationships, in which one party has knowledge of and is able to authenticate the other party, but the other party has to rely on mediated trust regarding the first party. This is common with banks, for example, where to access your account online you need to present the credentials you've established directly with the bank, while the bank authenticates itself to you using mediated authentication (an X.509 certificate issued by a commercial CA and presented using the TLS protocol).

We believe that in this case, trust establishment and bootstrapping, trust auditing, and trust revocation are normal trust lifecycle activities.

Where problems seem to arise is in those cases where two entities without an existing relationship attempt to determine whether or not, and to what extent, they may trust each other. With some exceptions (for example, unauthenticated IPsec SAs [RFC5386], or the use of self-signed X.509 certificates), this involves the use of some mediating agent and tends to rely on a transitive trust model.

Transitivity in trust is similar to transitivity in mathematics: "Things which are equal to the same thing are equal to one another." (the first of Euclid's Common Notions). When there are transitive trust relationships, if A trusts B and B trusts C, then A trusts C.

Quite possibly the most common case of mediated trust on the internet

is the use of X.509 [RFC5280]certificates in TLS [RFC5246]. X.509 certificates are issued to identify entities, but because of conflation of identity with trust issues are often seen as conveying trust. That is to say, a model in which I trust a given CA to assert identity is, in practice, often a seen as a model in which I trust a given entity based on its CA's assertion of identity.

A given user cannot reasonably be expected to have pre-installed end entity certificates for every server she is likely to want to access, and so we have mediated trust based on someone (in this case, a certification authority) making an authoritative statement about the identity of a server, and that statement being verifiable using formal and well-understood (??) validation procedures, walking a chain of trust back to an installed trust anchor.

Another example, but one in which communicating entities may be closely related and still not have foreknowledge of one another, is in the use of group keys, as in GDOI [RFC6407] (the Group DOI for IKE). In that case group members share a key, but access to the key (along with key management operations including initial authentication) is mediated through a Group Controller/Key Server.

A non-cryptographic example of mediated, transitive trust is in VoIP systems in which a call control server is used. For example, if user Customer A has service with Service A, and is able to authenticate to that service, and Customer B has service with Service B, and is similarly able to authenticate to its service, Customer A is able to talk to Customer B if Service A and Service B know about each other and trust each other.

A variation on the mediated, authority-based trust models described above is consensus systems, where an endpoint or user still needs to rely on an external source for the basis for trust decisions, but a trust decision is based on agreement (or not) between a number of parties. If a very large number of parties state that a given entity is trustworthy, with little disagreement, that leads to a different decision from one when there's substantial agreement that a given entity is untrustworthy, or when there's very little agreement (or insufficient data). As of this writing we are not familiar with consensus-based trust models in IETF protocols.

5. Problems

We believe that these are the major problems with the internet trust infrastructure as commonly used in IETF protocols:

- Users, services, and other network elements are often required to make trust decisions about entities with which they have no previous relationship

- There is often insufficient information about the practices at and reliability of network entities making identity and attribute assertions

- There has been no delegation mechanism to make it clear when one entity is authorized to act on behalf of another entity. OAuth is [3] an authorization mechanism currently under development which may prove to be useful for generalized service delegation.

As described in Section 4, we believe that where there are problems related to trust in IETF protocols, it is largely in situations in which participating endpoints have no foreknowledge of one another, or the knowledge is unilateral.

This is due at least in part to the familiar problem of conflating authentication - proven identity - with the problem of authorization. In the case where two entities have an existing relationship this is probably reasonable. It is unlikely that the credentials or resources would have been provisioned if the relationship were not authorized.

In cases where there is no pre-existing relationship, however, there is frequently insufficient basis to make a trust decision. Transitivity is not appropriate in all cases, and is a genuinely bad idea in many.

When a certification authority issues a certificate and signs it, they are making an identity assertion. That may be sufficient for access control decisions when there is local knowledge of the identity being asserted, or when the resources being requested are low-value or not sensitive. The broader problem with identity assertions is that it is not always possible to know how reliable, or trustworthy, a given certification authority or trust anchor is, or what their vetting practices are for verifying a customer's actual identity before issuing a certificate or other credentials. Having a certification authority vouch for an entity's identity is meaningless if the CA is not careful about making sure that an applicant really is who they say they are. Unfortunately it is not often possible to know how reliable a given CA is, or whether or not their vetting

practices meet a given set of requirements.

In addition to the limitations with the existing internet trust and identity infrastructure, there are some missing components, as well. There isn't always a mechanism to identify the relationship between two entities when one is needed. For example, a utility company (gas, electric, sewer, water) may use a third-party payments company, and when you use the utility's website, when you click the "Pay my bill" button you're taken to the payment company's website. From the underlying identity and trust mechanism it is not possible to determine that there really is a relationship between the utility and the payment company, and that the payment company is authorized to collect money on behalf of the utility. A delegation mechanism is missing.

6. Security Considerations

This document attempts to describe and identify problem areas related to trust in the internet infrastructure, within IETF scope. As such it does not introduce new mechanisms. However, it should be understood that the problems described in this document do have immediate impact on the security of related mechanisms. Recommendations for remediation are outside the scope of this document.

7. Path forward

We suggest that there may be value in pursuing discussion of some of the question raised earlier in this document. In particular,

- Is there value in a shared understanding or shared definition of what is meant by the word "trust?"
- Do we, as an organization, care about clearer descriptions of trust models in IETF protocols?
- Do we need to develop a stronger understanding of how to support trust frameworks, or how to develop frameworks in which multiple trust and policy models are used in a given scenario (say, in VoIP, where you may involved DNS, SIP, TLS, STUN, and others in completing a single "call")?
- Should we be differentiating between threats introduced by cryptographic or protocol flaws, and threats tied to trust problems?
- Does renewed interest in federated and other third-party identity and authorization mechanisms affect organizational priorities around trust issues?
- What, if anything, does the IETF need to be doing more generally around questions of trust?

8. Informative References

[tbss] Aljazzaf, Z., "Trust-Based Service Selection", December 2011, <http://ir.lib.uwo.ca/cgi/viewcontent.cgi?article=1478&context=etd>.

[ws-trust] "WS-Trust 1.3: OASIS Standard", March 2007, <http://docs.oasis-open.org/ws-sx/ws-trust/200512/ws-trust-1.3-os.html>.

[samlgloss] "Glossary for the OASIS Security Assertion Markup Language (SAML) V2.0", March 2005, <https://www.oasis-open.org/committees/download.php/21111/saml-glossary-2.0-os.html>.

[RFC4949] Shirey, R., "Internet Security Glossary, Version 2", RFC 4949, August 2007.

[RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008.

[RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, May 2008.

[RFC5386] Williams, N. and M. Richardson, "Better-Than-Nothing Security: An Unauthenticated Mode of IPsec", RFC 5386, November 2008.

[RFC5914] Housley, R., Ashmore, S., and C. Wallace, "Trust Anchor Format", RFC 5914, June 2010.

[RFC5934] Housley, R., Ashmore, S., and C. Wallace, "Trust Anchor Management Protocol (TAMP)", RFC 5934, August 2010.

[RFC6024] Reddy, R. and C. Wallace, "Trust Anchor Management Requirements", RFC 6024, October 2010.

[RFC6407] Weis, B., Rowles, S., and T. Hardjono, "The Group Domain of Interpretation", RFC 6407, October 2011.

[1] <http://www.identitywoman.net/the-trouble-with-trust-the-case-for-accountability-frameworks>

[2] <http://www.nist.gov/nstic/>

[3] <http://oauth.net/>

Authors' Addresses

Melinda Shore
No Mountain Software
PO Box 16271
Two Rivers, AK 99716
US

Phone: +1 907 322 9522
Email: melinda.shore@nomountain.net

Karen O'Donoghue
ISOC
7167 Goby Lane
King George, VA
US

Email: odonoghue@isoc.org