A Field Guide to Internet Trust

by Kaliya "Identity Woman" Hamlin and Steve Greenberg

Summary

The decreasing cost of computation and communication has made it easier than ever before to be a service provider, and has also made those services available to a broader range of consumers. New services are being created faster than anyone can manage or even track, and new devices are being connected at a blistering rate.

In order to manage the complexity, we need to be able to delegate the decisions to trustable systesms. We need specialists to write the rules for their own areas and auditors to verify that the rules are being followed.

This paper describes some of the common patterns in internet trust and discuss some of the ways that they point to an interoperable future where people are in greater control of their data. Each model offers a distinct set of advantages and disadvantages, and choosing the appropriate one will help you manage risk while providing the most services.

For each, we use a few, broad questions to focus the discussion:

- How easy is it for new participants to join? (Internet Scale)
- What mechanisms does this system use to manage risk? (Security)
- How much information the participants require from one another how strongly verified? (Level of Assurance -not what I think assurance is...but we can talk - it often also refers to the strength of security like number of factors of authentication)

Using the "T" Word

Like "*privacy*", "*security*", or "*love*", the words "*trust"* and "*identity"*, and *"scale"* carry so much meaning that any useful discussion has to begin with a note about how we're using the words.

This lets each link the others to past behavior and, hopefully, predict future actions. The very notion of trust acknowledges that there is some risk in any transaction (if there's no risk, I don't need to trust you) and we define *trust* roughly as:

> The willingness to allow someone else to make decisions on your behalf, based on the belief that your interests will not be harmed.

The requester trusts that the service provider will fulfill their request. The service provider trusts that the user won't abuse their privileges, or will pay some agreed amount for the service. Given this limited definition, *identity* allows the actors to place one another into context.

Trust is contextual. Doctors routinely decide on behalf of their patients that the benefits of some medication outweigh the potential side effects, or even that some part of their body should be removed. These activities could be extremely risky for the patient, and require confidence in the decisions of both the individual doctor and the overall system of medicine and science. That trust doesn't cross contexts to other risky activities. Permission to prescribe medication doesn't also grant doctors the ability to fly a passenger airplane or operate a nuclear reactor.

Trust is directional. Each party's trust decisions are independent, and are grounded in the identities that they provide to one another.

Trust is not symmetric. For example, a patient who allows a doctor to remove part of their body should not expect to be able to remove parts of the doctor's body in return. To the contrary, a patient who attempts to act in this way would likely face legal sanction.

Internet Scale

Services and APIs change faster than anyone can manage or even track. Dealing with this pace of change requires a new set of strategies and tools.

The general use of the term "Internet Scale" means the ability to process a high volume of transactions. This is an important consideration, but we believe that there is another aspect to consider. The global, distributed nature of the internet means that scale must also include the ease with which the system can absorb new participants. Can a participant join by clicking "Accept", or must they negotiate a custom agreement?

In order to make this new world of user controlled data possible, we must move from a model broad, monolithic agreements to smaller, specialized agreements that integrate with one another and can be updated independently.

A Tour of the Trust Models

The most straightforward identity model, the **sole source**, is best suited for environments where the data is very valuable or it is technically difficult for service providers to communicate with one another. In this situation, a service provider issues identity credentials to everyone it interacts with and does not recognize identities issued by anyone else. Enterprises employing employees, financial institutions, medical providers, and professional certifying organizations are commonly sole sources. Because this is the most straightforward model to implement, it is also the most common.

Two sole sources might decide that it's worthwhile to allow their users to exchange information with one another. In order to do so, they negotiate a specific agreement that covers only the two of them. This is called a **Pairwise Agreement** and, while it allows the two parties to access confidential resources, the need for a custom agreement makes it difficult to scale the number

of participants. This is also a kind of *federated identity model*, which simply means that a service accepts an identity that is managed someplace else.

As communication technology became more broadly available, the number of institutions who wanted to communicate with one another also increased. Groups of similar organizations still wanted to issue their own identities, but wanted their users to be able to interact freely with one another. The prospect of each service having to negotiate a custom agreement with every other service was daunting, so similarly chartered institutions came up with standard contracts that allow any two members to interact. These groups are called **Federations**, and there are several different kinds. Federation agreements and membership are managed by a **Contract Hub**.

When the federation agreement limits itself to policy, governance, and common roles, but leaves technical decisions to the individual members, it's referred to as a **Mesh Federations**. Individual members communicate form a mesh, and can communicate directly with one another using whatever technology they prefer.

Alternatively, a **Technical Federation** defines communication methods and protocols, but leaves specific governance and policy agreements to the members. In some cases, the technical federation may also route messages between the members.

As the number of services has increased, so has the problem of managing all of those usernames and passwords. Users might decide to reuse an existing identity rather than creating a new one. In recent years, some organizations have made identities that they issue available to other services. Service providers accept these identities because it lowers the cost of user acquisition. When the same entity provides identities for both the requester and the service provider, it is referred to as a **Three Party Model**.

If the requester and the service provider have provider have separate but compatible identity providers, it is called a **Four Party** model. This is present in highly dynamic models, such as credit card processing,

Peer-to-peer networks are for independent entities who want to identity assurance, but who lack a central service that can issue identities to everyone. To get around this, the participants vouch for one another's identities.

Individual contract wrappers are an innovation to enable complex connections between services where the terms and conditions of using the data are linked to the data.

Common Internet Trust Models

Sole source	A service provider only trusts identities that it has issued.
Pairwise Federation	Two organizations negotiate a specific agreement to trust identities issued by one another.
Peer-to-Peer	In the absence of any broader agreement, individuals authenticate and trust one another.
Three-Party Model	A common third party provides identities to both the requester and the service provider so that they can trust one another.
"Good Enough" Portable Identity	In the absence of any institutional agreement, service providers accept individual, user-asserted identities.
Federations	A single, standard contract defines a limited set of roles and technologies, allowing *similar* types of institution to trust identities issued by one another.
Four-Party Model	An interlocking, comprehensive *set* of contracts allows different types of entity to trust one another for particular types of transaction.
Centralized Token Issuance, Distributed Enrollment	A shared, central authority issues a high-trust communication token. Each service provider independently verifies and authorizes the identity, but trusts the token to authenticate messages.
Individual Contract Wrappers	Manage how personal data is used rather than trying to control collection. Information is paired contract terms that governs how it can be used. Compliance is held accountable using contract law.
Open Trust Framework Listing	An open marketplace for listing diverse trust frameworks and approved assessors.
Personal Cloud + Agents	An Individual has a personal Cloud and delegates agents it trust to work on their behalf.

Peer-to-Peer Identity

There is no central identity provider in a peer-to-peer trust network. Users assert their own identities and each individual decides who they trust and which they do not. This can provide a high level of trust once identities have been established, but technological complexity has hindered adoption.

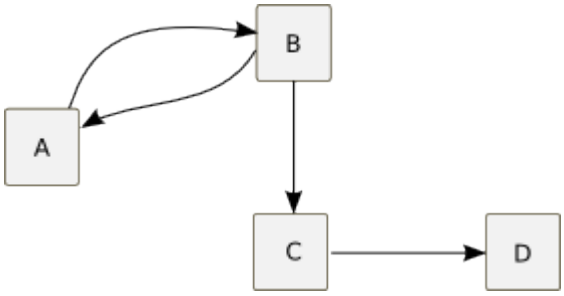

Advantages: No dependence on a central identity provider. Users can assert any identity that they want. A high degree of trust can be achieved, once identities have been exchanged and authenticated.

Disadvantages: Requires a high degree of technical sophistication. Individually verifying each entity can be labor intensive. Tracking identities that have been revoked can be complex and error prone.

Ability to Scale: Very difficult to scale to a large number of users.

Examples: The most successful peer-to-peer identity networks on the Internet have been implemented using *public key* cryptography, which allows participants to trust messages sent over insecure channels. Peer-to-peer identities are used to securely exchange email and instant messages, and to verify the integrity of software packages. [should we name PGP?]

Note that the use of public key techniques does not necessarily mean that a peer-to-peer identity model is in use. Other applications - such as the SSL/TLS system used to encrypt web traffic - use public keys to verify that a trusted third party has authenticated an identity

Sole Source

A service provider issues identities to all users and only trusts identities that it has issued. This model is also referred to as an “Identity Island”, because it is not connected to anything else. The service provider performs its own verification and usually dictates governance, privacy, and technical terms to the requester. There is minimal - if any - negotiation between the requester and the service provider.

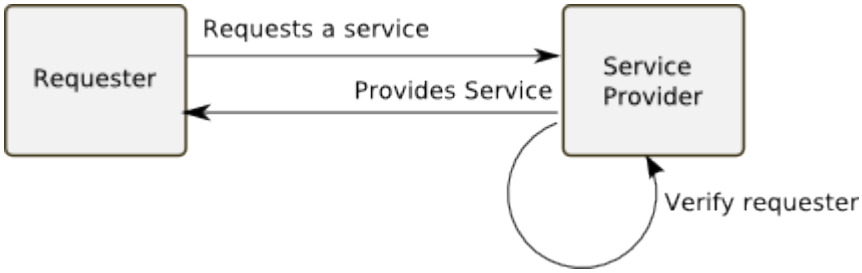

When to Use: A service that maintains particularly confidential information or valuable assets, or that operates in an uncertain environment. If proper operation and risk management requires a high level of assurance, then consider being a sole source.

Advantages: The service provider can authenticate requesters to whatever level of assurance it desires before issuing an identity and does not depend upon third parties.

Disadvantages: The service provider bears the full management cost of the identity life cycle. The requirement to create a new identity may discourage potential users of the service. The service must provide an attractive enough product to justify asking the requester to create and manage a new account.

Being a sole source provider does not guarantee account security, as end users may simply give their credentials (login and password)to a third party. Tricking users into giving up account information is a common tactic used by “phishing” sites and other criminals, but legitimate services like Mint.com (for managing the data about all ones financial accounts from various institutions) do this as well.

Ability To Scale: When the service provider does not need to integrate with any other services or when it is in a position to dictate terms, a sole source trust model can scale to very large systems. The requirement to create and remember new identity can be a barrier to growing the number of active users.

Examples: Historically, this has been the most common identity model because it can be implemented simply and gives the service provider the most control. Large, consumer-facing services like eBay, Facebook, and Yahoo! were created with sole source identity, although many are adopting newer models as internet technology has evolves. Internal corporate services are generally sole source, and only accept identities issued by the organization.

Financial services, health insurance, and government are likely to remain sole source identity providers until a strong, multifactor identity gains momentum with consumers. There have been several attempts to do this, but none has yet achieved critical mass.

Pairwise Agreement

Two institutions want to trust identities issued by one another, but there is no existing way for them to do so. They negotiate a specific agreement that covers only the two of them. A pairwise agreement might specify governance, security and verification policies, or specific technical methods.

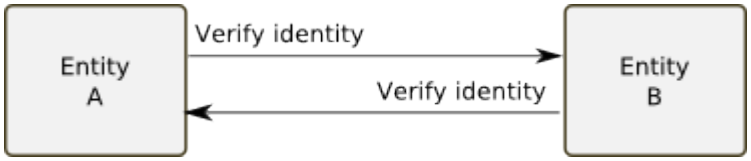

When to Use: Business or institutional partners want to grant one another access to confidential systems or information, but no standard contracts or umbrella organizations exist.

Advantages: Organizations can grant one another access to scarce resources and confidential information. Highly customized for the specific situation and participants.

Disadvantages: Time consuming and complex to negotiate, expensive. Difficult to scale.

Ability to Scale: Pairwise federations do not scale well, because each additional party will need to make a custom agreement with every other party.

Examples: Businesses might negotiate pairwise agreements with large supplier. Educational institutions may craft specific research agreements.

One of the main products from Liberty Alliance was the emergence of technical (SAML 1 and 2) and legal standards for pairwise federation. There are 100,000's of these in operation today.