A Field Guide to Internet Trust

by Kaliya "Identity Woman" Hamlin and Steve Greenberg
Contact: Kaliya [@] Identitywoman [dot] net
Steve [at] sixgables [dot] org

Introduction

In this paper we describe several models of federated identity in common use on the Internet and propose names for them. Our goal is to propose a common vocabulary that can simplify and clarify discussions about identity and trust models on the Internet. We begin by offering limited definitions for the words "identity", "trust", and "internet scale" and then use those to show how the models differ from one another.

The "T" Word and Friends

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

Identity

For the purpose of this conversation, an *Identity* is a recognizable token allows us to reliably tell one entity from another, and that does so consistently through time. Identity lets us link a request to past behavior and helps us predict future actions. In this way, it places the request into a context. Every interaction carries some risk, and identities provide a foundation for assessing and managing that risk. When we don't know one another directly and depend upon a third party to authenticate us, identities are said to be *federated*. Each of the models laid out in this paper reflect different decisions about who can provide identity and how it should be used.

Trust

Having identified each other, the next question becomes what we allow one another to do. A user may be asking for access to a resource, or a service provider might store some piece of data. The requester trusts that the service provider will answer accurately and in a timely manner. The service provider trusts that the user won't abuse their privileges, or will pay some agreed amount for the service. Stated generally, we define trust as:

> *The willingness to allow someone else to make decisions on your behalf, based on the belief that your interests will not be harmed.* [1]

Several factors influence the choice about when to allow someone else to make a decision, and to what extent:

- **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, however. Permission to

[1] Our model of trust draws heavily on the work of Aaron Hoffman, particularly "*A Conceptualization of Trust in International Relations*" - European Journal of International Relations, Vol 8, 2002.

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

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 distributed nature of the internet means that the notion of scale must also include the ease with which the network can absorb new participants. Must a new participant negotiate a custom agreement with every other member, can they sign a standard contract, or is there no contract at all? In the context of federated identity we consider how easy it is for service providers to join the network and accept identities, and also how easily new requester identities can be created for use with services.

Common Internet Trust Models

Sole source	A service provider only trusts identities that it has issued.
Pairwise Agreement	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.
“Bring Your Own” Portable Identity	The service provider specifies the technical methods that it will accept and the requester can choose any identity provider they like.
Federations	Pre-negotiated standard contracts defines roles and governance so that organizations can trust one another without having to negotiate directly.
Four-Party Model	An interlocking, comprehensive *set* of contracts allow risk assessors to specialize.
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	Information is paired with contract terms that govern how it can be used. Compliance is held accountable using contract law.
Open Trust Framework	An open marketplace for listing diverse trust frameworks and approved assessors.

Sole Source

A single service issues identities and only trusts identities that it has issued. This model does not federate identities at all, and is sometimes referred to as an “Identity Island” because it is not connected to anything else. The service provider performs its own verification and dictates governance, privacy, and technical terms to all participants.

There is minimal - if any - negotiation between the requester and the service provider. The service provider manages the entire account lifecycle from creation through retirement.

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 often sole source, and only accept identities issued by the organization.

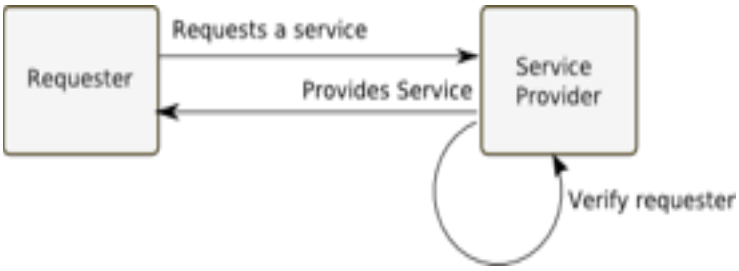

Financial services, and health insurance, are likely to remain sole source identity providers until a strong, multifactor identity gains momentum with consumers and liability questions are settled. There have been several attempts to do this, but none has yet achieved critical mass.

Being a sole source provider does not guarantee account security, as end users may simply give their account 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 (a US-based financial service provider) also ask for credentials in order to combine information from sites that do not provide APIs.

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 a product attractive enough to justify asking the requester to create and manage a new account.

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.

Peer-to-Peer Identity

When no central identity provider or governance agreement is present, participants assert their own identities and each individual decides who they trust and who they do not. Each participant is a peer with equal standing and each can communicate with anyone else in the network.

Examples: The most familiar peer-to-peer network is probably e-mail. An internet host can join the e-mail network with little more effort than updating its DNS entry and installing some software. Once a host has joined the network, individual e-mail addresses are easily created with no requirement for approval by any central authority. This flexibility and ease of account creation helped spur the growth of the internet, but also allows spam marketers to create false emails.

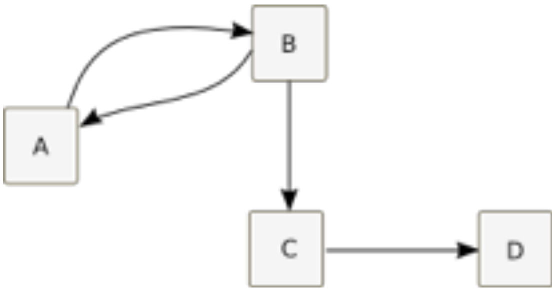

The best known secure peer-to-peer identity networks on the Internet have been implemented using *public key cryptography*, which allows participants to trust messages sent over insecure channels like email. Products like PGP and it's open source counterpart gpg are the most common implementations of public key messaging tools.

When To Use: No central identity provider is available but network participants can exchange credentials.

Advantages: No dependence on a central identity provider. No formal agreement needed to join the network. Participants can assert any identity that they want. Secure peer-to-peer technologies can provide a high degree of confidence once identities have been exchanged. Peer-to-peer models are very flexible, and can support a wide range of trust policies.

Disadvantages: No governing agreement or requirement to implement any policies. Secure deployment requires a high degree of technical sophistication and active management. Individually verifying each participant can be labor intensive. Tracking identities that have been revoked can be complex and error prone.

Ability to Scale: If security requirements are low, peer-to-peer networks can grow very large because new members can join easily. Higher levels of security can be complex to deploy and operate, and can impose a practical limit on the size of the network.

Centralized Token Issuance, Distributed Enrollment

A special case peer-to-peer network. Participants want to establish trusted identities that can be used securely for ongoing, high-value communication among organizations. A trusted, central provider issues identity tokens which are then enrolled independently by each service provider. Service providers are not required to cooperate or accept one another's enrollments.

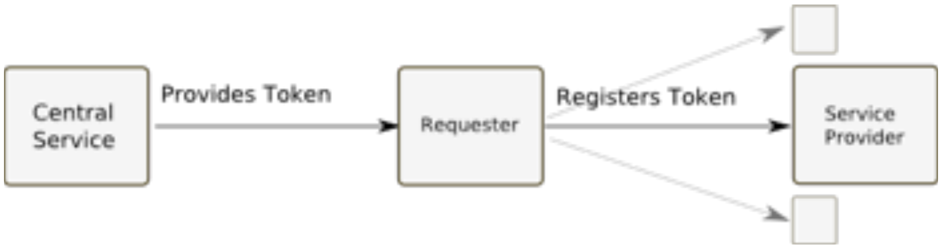

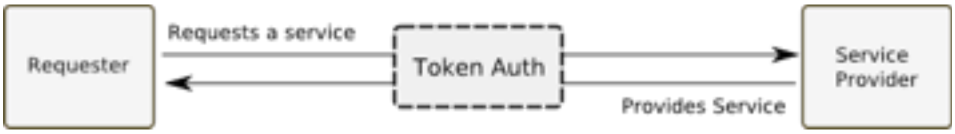

Examples: The most common examples are RSA SecurID and SWIFT 3SKey. Hardware tokens are issued by a trusted provider, which are then used to authenticate individual identities.

Each service will require the user to enroll separately, but once the user has registered they can use the token for future interactions.

When the requester wants to use a service, they're authenticated using the token.

When to use: Strong Authentication across a range of business entities who may have different enrollment requirements.

Advantages: Can provide a high level of identity assurance to institutions spread across legal and national boundaries.

Disadvantages: Can be expensive and complex to implement. Depends upon the existence of a trusted third party who can issue and ensure the security of hardware tokens. Hardware tokens can be lost.

Ability to scale: Can scale to large networks.

Pairwise Agreement

Two institutions want to trust identities issued by one another, but there is no outside governance or policy framework for them to do so. They negotiate a specific agreement that covers only the two of them. Each institution trusts the other to properly manage the identities that it issues.

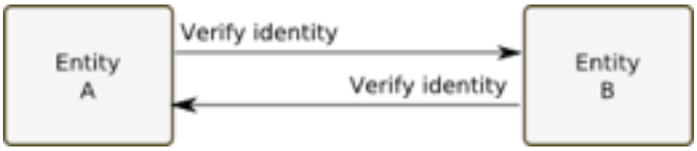

Examples: A pairwise agreement can specify governance, security and verification policies, or specific technical methods.

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

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.

Three Party Model

A trusted third party provides identities to both the requester and service provider. In order to interact with one another, both must agree to trust *the same* identity provider.

Examples: Google, Facebook, American Express, Paypal, Amazon, iTunes App Store

There are two broad types of Three Party Model. If one (or both) of the parties insists on a particular identity provider, we refer to it as a *Winner Take All* network because other identity providers are locked out. If only technical methods are specified and the requester is free to specify any identity provider they like, we refer to it as a *Bring Your Own Identity* network.

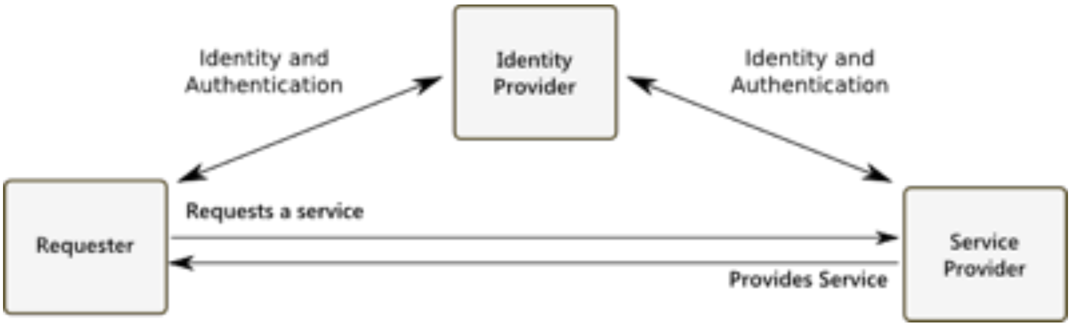

When to Use: An identity provider may choose to offer a three party model when it can provide identities more efficiently than the requester or service provider can on their own. Requesters and service providers may choose to implement a three party network for access to an existing market.

Advantages: Separates identity management from the service being provided. In cases where a shared third party is available, this model simplifies the process of exchanging trusted identities. Malicious actors can be identified and isolated from the entire network. Requesters can use a single identity with many service providers, and service providers can trust requesters without having to verify each one.

Disadvantages: Because participants can only interact if they have been authenticated by a single identity provider, that provider wields substantial power. The identity provider effectively controls the requester's ability to use services and the services' ability to work with requesters.

For instance, a requester who loses their account with the identity provider also loses all of the services where they used that identity. If you use your Facebook to sign in to other products then you also lose those other products if your Facebook account is closed.

Ability to Scale: Very difficult to get started because a three party network is not interesting to service providers until it has users, but only attracts users if it has interesting services. Once they are established and functioning, however, a successful three party network can grow extremely large.

“Bring your Own Identity” Three Party Model

A special case of the three party model where the service provider specifies the technical methods that it will accept, but allows the requester to choose any identity service they like. The service provider does not set details for identity verification or authentication and simply assumes that the requester has chosen one that's good enough for their purposes. The service provider and requester agree to terms, the requester and the identity provider agree to terms, but the service provider does not make any agreement with the identity provider.

Examples: The most common Bring Your Own Identity technologies are SAML, OpenID, and email address verification.

When to Use: The service provider does not want to bear the cost of managing the requester's identity, or wants to simplify account creation and sign-in.

Advantages: The requester can use an existing identity rather than having to create a new one for this service. If the requester chooses a good identity provider, the service gets the benefit of higher security with no additional cost.

Disadvantages: The account is only as secure as the authenticating service. The service provider depends on the user to select a trustworthy identity service.

Designing a user interface that allows the user to specify an identity provider has proved to be difficult. Consumers don't generally have the experience to know a good identity provider from a bad one so, in practice, they depend upon seeing a familiar brand. When OpenID was first introduced, supporting sites attempted to help by listing a large set of brands so that the user could choose a familiar one. The resulting products ended up so festooned with logos that they were likened to NASCAR cars, and ended up being more confusing than helpful.

Ability to Scale: Very high.

"Winner Take All" Three Party Model

A special case of the three party model where the service provider wants to allow the requester to use an existing identity, but only accepts authentication from a defined set of providers. Participants sign an agreement with the identity provider, which also allows them to talk to one another.

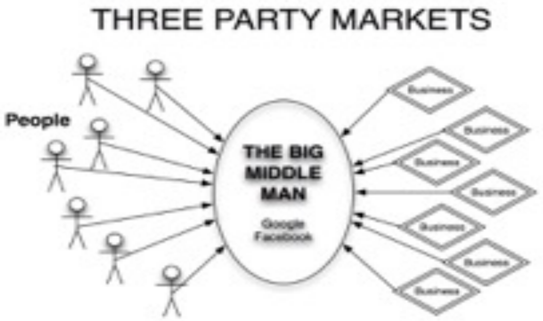

Examples: Apple completely controls the channel between app vendors and iPhone users, deciding which applications are available and which users are allowed to use them. Spotify and Zynga games depend upon Facebook for authentication.

When to Use: The service provider wants to take part in a large, established channel, or requires a high level of assurance.

Advantages: The requester can use an existing identity, which lowers the amount of effort required to use a new service. The service provider gets access to the users of an identity network without having to manage the accounts itself. Some identity providers offer higher security than the service could practically provide on its own.

Large three-party model identity providers like Facebook, Google, and PayPal dedicate substantial resources to security.

Disadvantages: Because participants can only interact if they have been authenticated by a single identity provider, that provider wields substantial power. The identity provider effectively controls the requester's ability to use other company's products. For instance, a requester who loses their account with the identity provider also loses all of the services where they used that identity. If you use your Facebook to sign in to other products then you also lose those other products if your Facebook account is closed.

Conversely, a service provider that depends on a single third party identity provider leaves themselves open to the third party deciding to change its terms.

Ability to Scale: Difficult to get started because it is only interesting to service providers when it has consumers, but only interesting to consumers if it can offer interesting services. Once they are established and functioning, however, a successful identity provider can build a very large network.

Federations

A *Federation* provides a standard, pre-negotiated set of contracts that allow organizations to recognize identities issued by one another. A federation agreement might specify user roles, governance, security and verification policies, or specific technical methods. The federation is organized around a *Contract Hub*, which is responsible for the agreements. Organizations with similar goals or structure create a standard agreement rather than negotiating individually.

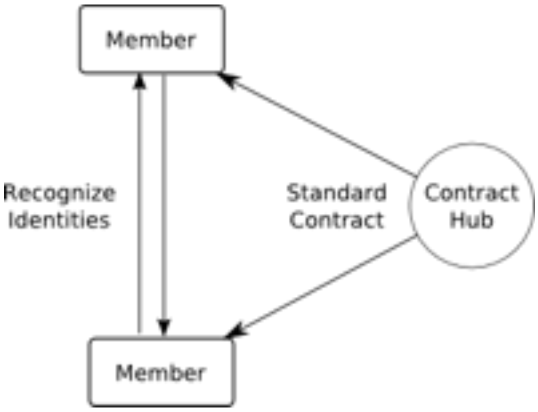

When to Use: A large number of organizations can agree upon roles and governance, and can create a standard contract.

Advantages: Organizations can recognize identities that one another issue without having to negotiate individual agreements with every party.

Disadvantages: Not customized for individual member organizations. Because of the need to create an agreement that a large number of parties can agree to, the federation might be limited to lowest common denominator roles.

Ability to Scale: Very high.

Mesh Federation

A *Mesh* Federation provides a legal and policy umbrella so that institutions can interact with one another but does not specify technical methods. Each member organization issues digital identities for its people and the federation agreement provides the legal framework for them to use one another's resources. The federation agreement might specify governance, policy, or roles, but the member institutions are free to implement using whatever technologies they like. This is referred to as a mesh because participating services connect directly with one one another in order to authenticate identities. For contrast, a federation network that provides a central identity clearing house is referred to a *Technical* federation (discussed below).

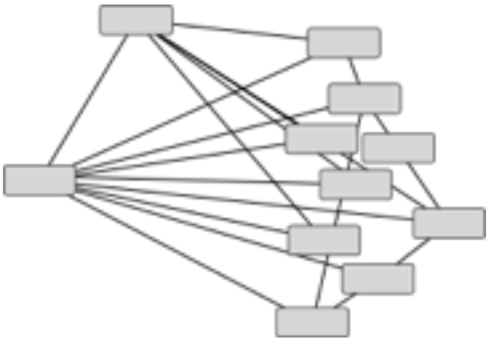

Examples: Mesh federations were pioneered by educational institutions. Universities already had a culture of cooperation and realized that the interest of students and research goals of faculty were best served by the free flow of information. NRENS (National Research and Education Networks) around the world include InCommon in the US, SurfNET in the Netherlands, and JISC/Janet in the UK.

When to use: Large institutions wish to share resources and can agree on roles and governance, but do not need a central point for authenticating identity.

Advantages: Federation participants don't need to negotiate custom agreements with every other member.

Disadvantages: Because of the need to gather broad adoption, mesh federations may be limited to the most common roles and might not cover complex use cases.

Ability to Scale: Because the mesh federation provides a standard contract, it scales to a large number of members.

Technical Federation

In addition to contract terms, a *Technical* federation also provides a central service that acts as a clearinghouse for identity operations. It routes authentication requests from the service back to the requester's chosen identity provider, translating protocols as needed. The existence of a central service lowers the technical and administrative costs of participating in the network.[2] For contrast, a federation network where the participants connect directly with one another rather than going through a central clearinghouse is called a *Mesh*.

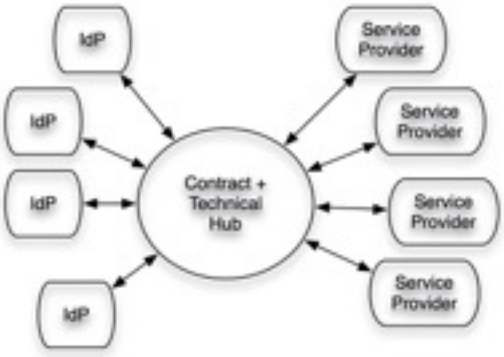

Examples: WAYF provides federated single sign-on to Denmark's higher education, research institutions, and libraries. [3]

When to Use: A large entity is available to act as an identity clearing house.

Advantages: Encourages use of digital identity by providing a central clearinghouse for authentication. Service providers only need to integrate with a single identity provider. Requesters can choose from a variety of identity providers.

Disadvantages: Requires substantial investment that may only be available to very large institutions or states.

Ability to Scale: Can scale to support national identity programs.

[2] http://www.wayf.dk/wayfweb/faq_-_frequently_asked_questions_attchmt/2008_11_23_tnc_abstract_trusted_third_party_based_id_federation_enhancing_privacy_and_lowering_the_bar_for_connecting%283%29.pdf

[3] http://www.wayf.dk/en/about-wayf

Inter-Federation Federations

When organizations are unable to communicate directly with one another because of legal limits or national boundaries, existing federations can negotiate *inter-federation* federations which allow members of different federations to interact with one another.

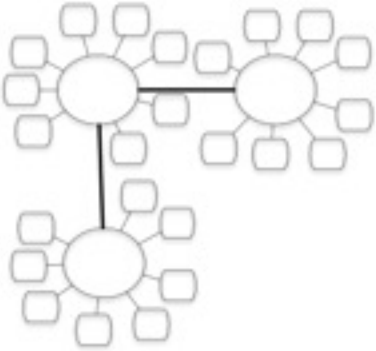

Examples: REFEDS[4], eduGAIN[5], and Kalmar2[6] are inter-federation programs for research institutions and higher education.

When to use: Institutions are unable to form direct relationships with one another because of legal or national boundaries, but have existing federations that can negotiate on their behalf.

Advantages: Federations can act as agents, negotiating for members to simplify the complexity of getting agreement among a large number of institutions.

Disadvantages: The complexity of negotiating inter-federation agreements slows the process and may limit the interactions that are covered.

[4] https://refeds.terena.org/index.php/Main_Page

[5] http://www.geant.net/service/eduGAIN/Pages/home.aspx

[6] https://www.kalmar2.org/kalmar2web/front_page.html

Four-Party Model

A *four-party* model provides a comprehensive set of interlocking legal contracts that detail roles, responsibilities, and technical methods. In order to take part in the network, each party must agree to one of the contracts in a given framework. Identity providers specialize in providing support for particular roles.

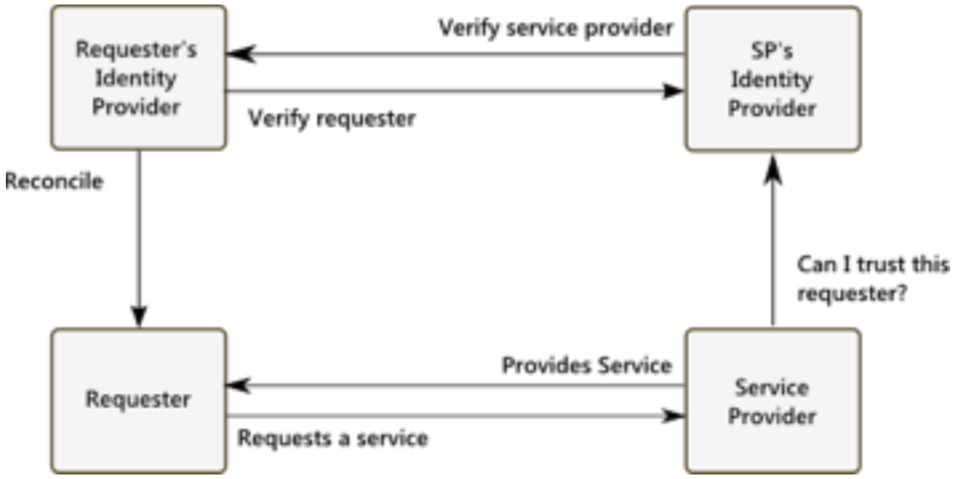

Examples: The credit card networks, such as Visa and Mastercard, are implemented as four party networks.[7] These represent a large collection of individuals and institutions, each of which must routinely trust participants they've never encountered before.

Parties of all types continually join and leave the network, making it impractical for any single organization to track them all. By creating a standard set of well defined roles that work together, the Visa and Mastercard enable risk assessors to specialize.

Because of the vast difference in the size of the entities involved (anywhere from an individual person to a multi-national corporation), and the complexity of governing law, no single contract could be both complete and understandable by all parties.

To solve this problem, the network created a comprehensive, interlocking set of contracts that lay out all of the roles that entities can play. For each role, the appropriate contract specifies the interactions and responsibilities. The network design allows for multiple identity providers, each of whom can specialize in managing risk for a particular set of users. Risks are managed at the system level.

When to use: Closed network where all parties can be expected to sign a contract to join.

Advantages: Enables a network where participants of different sizes can interact smoothly with one another.   Allows for specialization of risk management in a complex, constantly changing network where participants frequently join and leave.

Disadvantages: Depends upon the ability to create comprehensive contracts. Risk management can impose substantial costs on the network.

Ability to scale: Four party models can scale to a large number of participants.

[7] http://en.wikipedia.org/wiki/Card_schemes

Individual Contract Wrappers

When providing information to a service, the requester also provides terms for how that information can be used.[8][9] Service providers agree to honor those terms in exchange for access to the data, and compliance is enforced through contract law. Terms might include an expiration date, limits on whether the data can be re-sold, or whether it can be used in aggregate form. This model is the mirror image of the Sole Source.

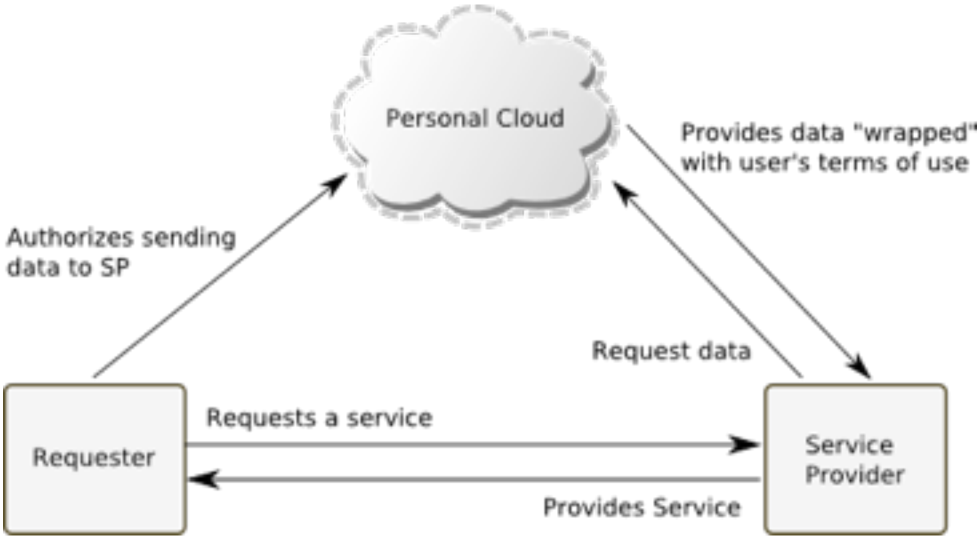

Examples: Personal.com offers a service that provides end users with a place to store personal data. Service providers agree to abide by a set of agreements in order to use this data.

When to use:

Advantages: Provides an incentive for the requester to provide clear, correct, and up-to-date information. In exchange for accepting limits on how the data can be used, the service provider gains access to better quality and more complete data.

Disadvantages: Emerging technology with evolving standards, not widely supported yet.

Ability to scale: It has a high ability to scale but it is almost a reverse architecture of the Sole Source and some of the same challenge.

[8] https://www.personal.com/legal-protection

[9] http://civics.com/rules-codificatio/

Open Trust Frameworks

A *Trust Framework* is a specification that describes a set of identity proofing, security, and privacy policies[10]. The framework is authored by subject matter experts, and is written with the intent that compliance can be assessed. The framework also lists the qualifications that an assessor must have in order to judge compliance.

A *Framework Listing Service* provides a publicly visible location where trust frameworks can be published and tracked. The listing service sets guidelines for acceptable frameworks and accredits assessors to verify that services implement the frameworks properly.

Examples: The Open Identity Exchange (OIX)[11], Kantara Initiative[12], and InCommon operate framework listing services. A Framework Creator authors a trust framework that specifies identity validation policies and publishes it to a Framework Listing Service. The framework may also specify the qualifications required in order to be a valid assessor of the policy.

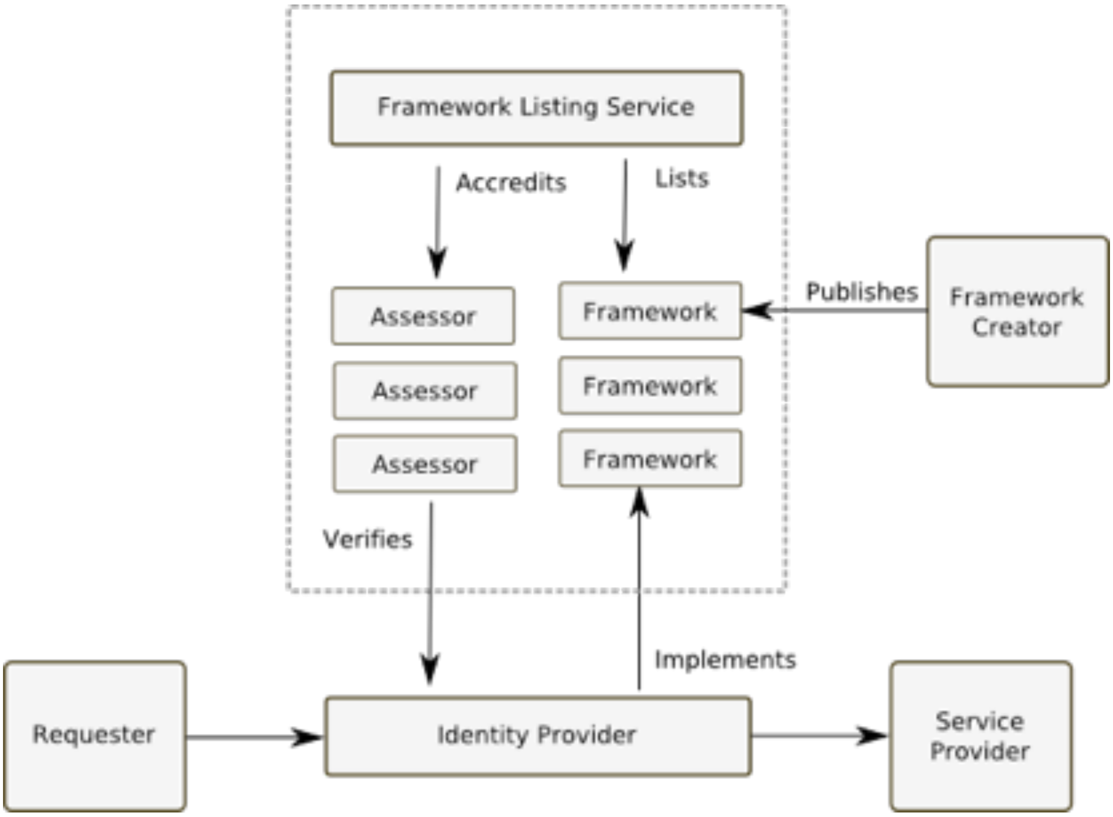

When to use: This should be used by networks who share a common set of technology and policy needs but are not in the business of creating technology networks or accrediting compliance.

Advantages: Standard, publicly available specifications that are designed by subject matter experts. Assessors can verify that the frameworks are implemented properly.

Disadvantages: Not broadly supported, evolving model.

Ability to scale: Because each component can be independently updated, a network based on open trust frameworks could potentially scale to be very large.

[10] http://openidentityexchange.org/what-is-a-trust-framework

[11] http://openidentityexchange.org/what-is-a-trust-framework

[12] http://kantarainitiative.org/confluence/display/certification/Identity+Assurance+Accreditation+and+Approval+Program