

Introducing the



TRUST
Over **IP**
FOUNDATION

On 5 May 2020, the Linux Foundation announced a new addition to its roster of global open source ecosystem projects: the **Trust over IP Foundation**. The mission of this new Foundation is to simplify and standardize how trust is established online so that everyone can feel safe, secure, and private in all of our digital interactions—whether between individuals, businesses, governments, or any “thing” on the Internet of Things.

In this white paper, we will cover:

- **Trust in the Pre-Internet Era**—the simple, global mechanisms we evolved to establish trust in relationships before we ever went online.
- **The Internet Era and the “Trust Gap”**—What happened when we moved online and why we ended out with such a large “trust gap” vs. real-world trust.
- **The New Era of Digital Trust**—How we can finally bridge this trust gap with open standard digital credentials and governance frameworks.
- **The Trust over IP Stack**¹—How this four-layer, dual-stack architecture has the potential to do for the peer-to-peer exchange of trustworthy digital credentials what the TCP/IP stack did for the peer-to-peer exchange of data packets.
- **The Role of the Trust over IP Foundation**—How this new organization will provide a global forum for collaboration on developing, hardening, testing, and promoting the Trust over IP stack.

¹ The starting definition of the ToIP stack is published as Hyperledger Aries RFC 0289: <https://github.com/hyperledger/aries-rfcs/blob/master/concepts/0289-toip-stack/README.md>

Part One: Trust in the Pre-Internet Era

Credentials

In the era before digital networks—when relationships and business interactions were all managed face-to-face—we had evolved a simple, universal, decentralized mechanism for achieving trust. We used **credentials** of all kinds.

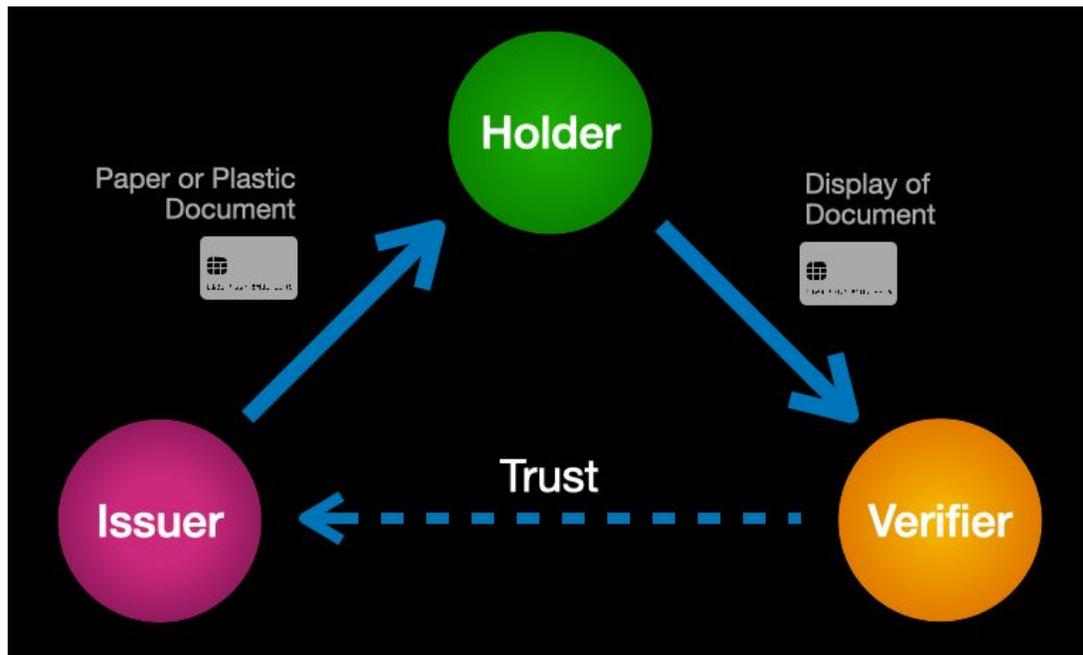


Note that by “credentials” we don’t just mean the pieces of paper or plastic that you carry around in your wallet to prove your identity, for example, driving licenses, government IDs, employment cards, credit cards, and so on. We mean any document of any size that enables you—or your organization—to prove something about you that enables the establishment of trust. For example, this could include:

- A birth certificate issued by a hospital or vital statistics agency that proves when and where you were born and who were your parents.
- A business registration or license of any kind that proves you are authorized to conduct a specific type of business.
- A diploma issued by a university that proves you have an educational degree.
- A passport issued by a government of a country that proves you are a citizen.
- An official pilot’s license that proves you can fly a plane.
- A utility bill that proves you are a registered customer of the utility.
- A power of attorney issued by the appropriate authority within a jurisdiction that proves that you can legally perform certain actions on behalf of another person.

The Credential Trust Triangle

The reason credentials have evolved as a universal mechanism for establishing real-world trust is the fundamental “trust triangle” illustrated below.

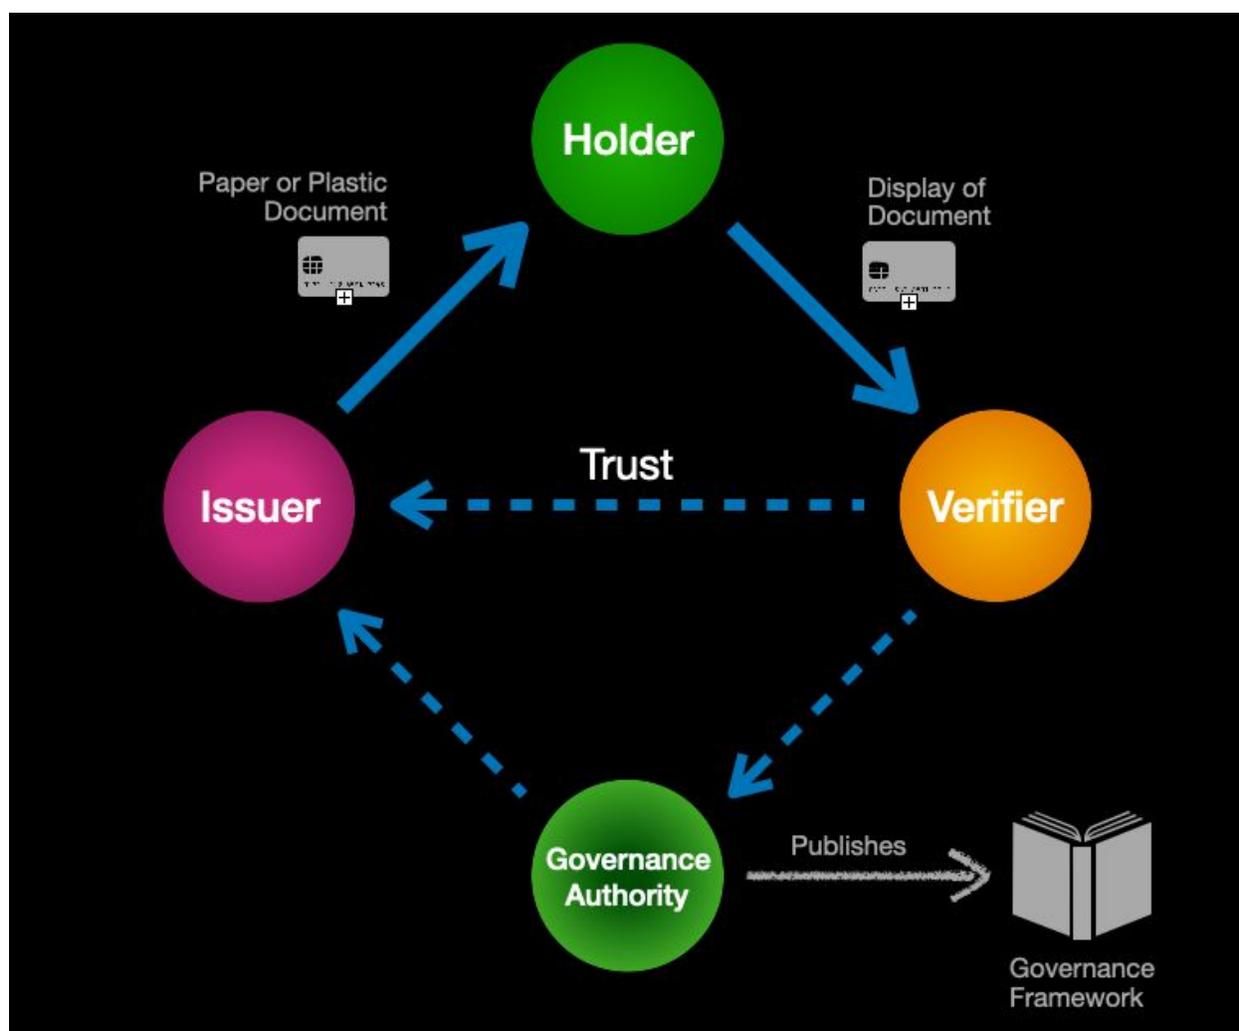


No matter what type of credential, the triangle involves the same three primary roles:

1. **Issuers** are the source of credentials—every credential has an issuer. Most are organizations such as government agencies (passports), financial institutions (credit cards), universities (degrees), corporations (employment IDs), churches (awards), etc. However individuals can also be issuers.
2. **Holders** request credentials from issuers, hold them in their wallets or filing cabinets, and present them when requested by verifiers (and approved by the holder). Although we most commonly think of individuals as holders, holders can also be organizations, or even things (such as the registration for a car).
3. **Verifiers** can be anyone seeking trust assurance of some kind about the holder of a credential. Verifiers request the credentials they need and then follow their own policy to verify their authenticity and validity. For example, a TSA agent at an airport will look for specific features of a passport or driver’s license to see if it is valid, then check to ensure it is not expired.

The Governance Trust Triangle

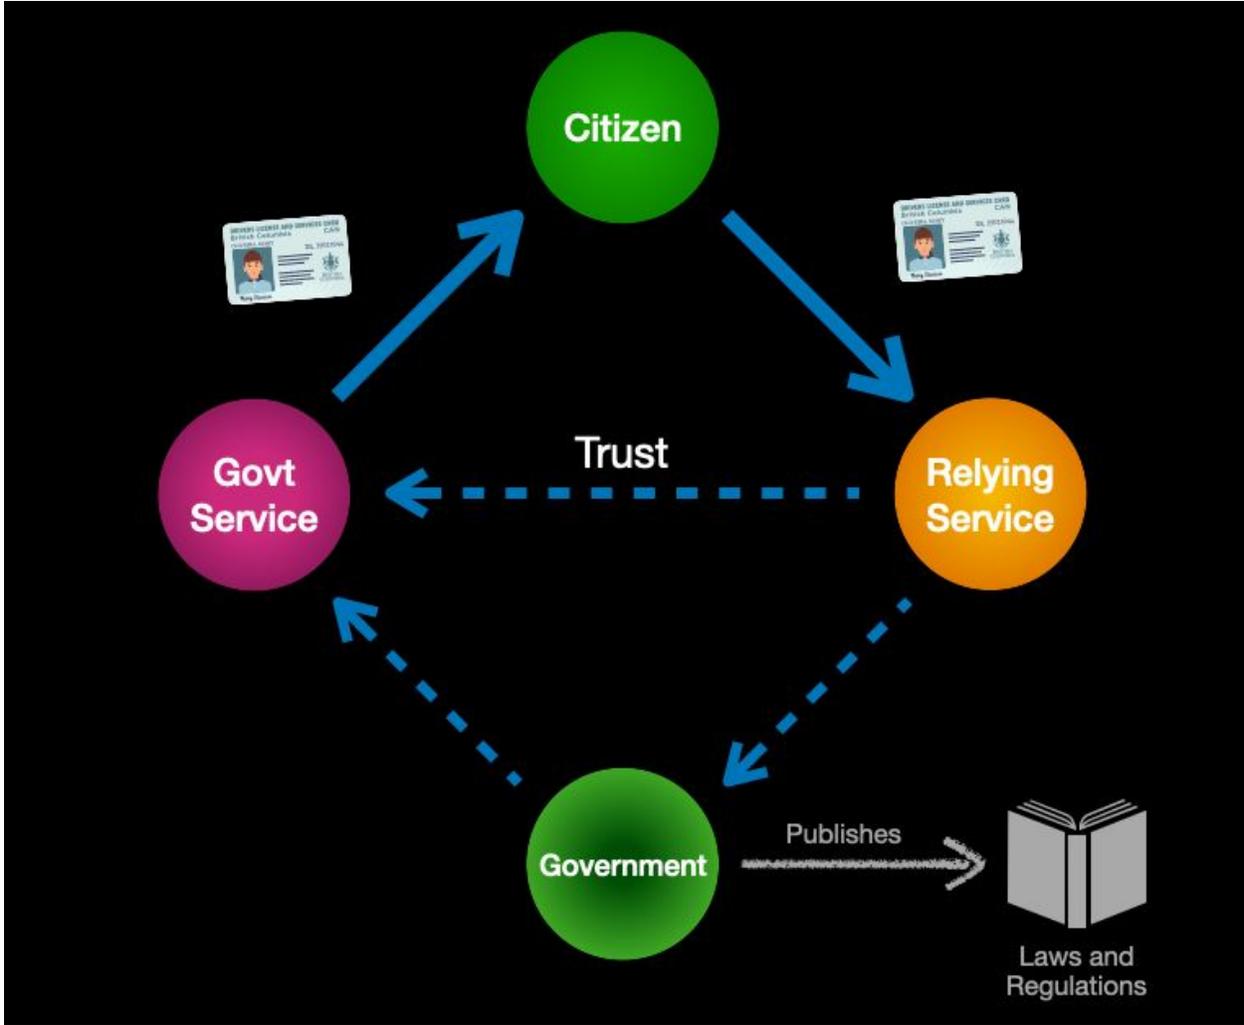
While some credentials only have a single issuer, others can be issued by many issuers. For example, passports are issued by hundreds of countries, and credit cards are issued by tens of thousands of banks and credit unions. For any credential that will be widely used by many holders and honored by many verifiers, there is a second trust triangle—the **governance trust triangle**—as shown below.



A governance authority can represent any set of issuers who want to standardize the business, legal, and technical policies for issuing, holding, and verifying a set of credentials. A governance authority can take any form—government, consortia, cooperative—but the purpose is the same: publish a **governance framework** that documents the rules by which the members of a trust community agreed to abide.

Example: Governmental Credentials

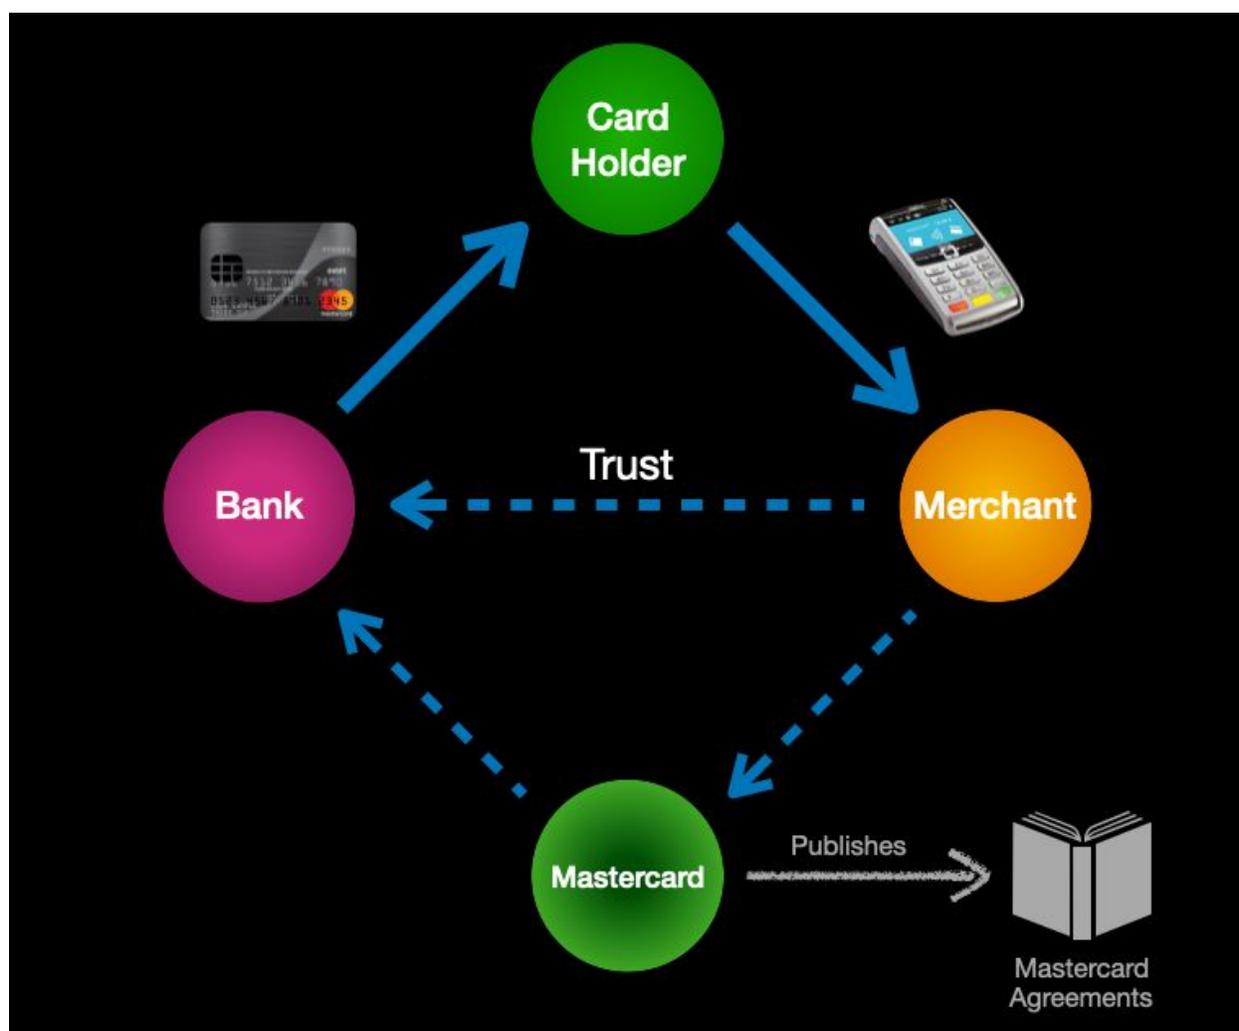
The best-known example of governance trust triangles are credentials issued by national governments following their own laws and regulations. Many (but not all) countries issue national citizen ID cards, and almost every country in the world issues passports under the [ISO/IEC 7810](https://www.iso.org/standard/50438.html) ID-3 standard.



In these governance trust triangles, the government itself is the governance authority, the laws and regulations of the country are the governance framework, and the issuers are the various government services authorized to issue a specific type of credential. Citizens or businesses can obtain the credentials for which they are qualified and then present them to any relying service that trusts the government for the accuracy of the information on the credential.

Example: Payment Card Networks

Another widely-known example of the governance trust triangle is a payment card network like Mastercard. In this case, the governance authority is Mastercard; the issuers are the banks and credit unions in the Mastercard network; the holders are the individuals or businesses that apply for Mastercards; and the verifiers are merchants enrolled in the Mastercard network to accept payment cards as shown below.

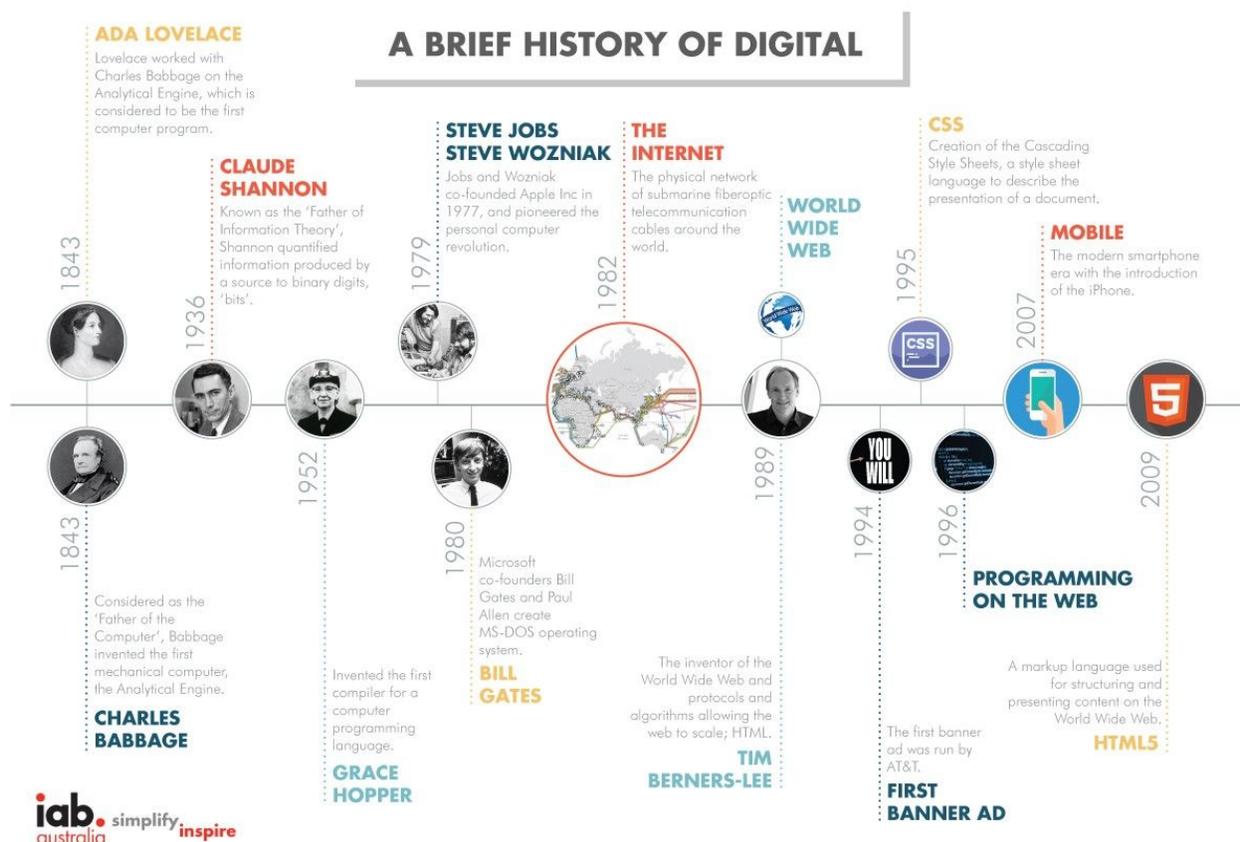


Although government-issued IDs and credit cards are the most common examples of credentials we carry in our own wallets, there are hundreds of other examples of governance trust triangles all around us: health insurance cards, student ID cards, employment ID cards, membership cards, loyalty cards, etc. In every case, the value of the credential depends on the trust the verifier has in the governance authority.

Part Two: The Internet Era and the “Trust Gap”

Moving Online

With the widespread consumer and business adoption of the Internet starting in the late 1990s, we moved into a new era of online relationships and digital interactions. The following [diagram from IAB Australia](#) puts this journey in perspective.



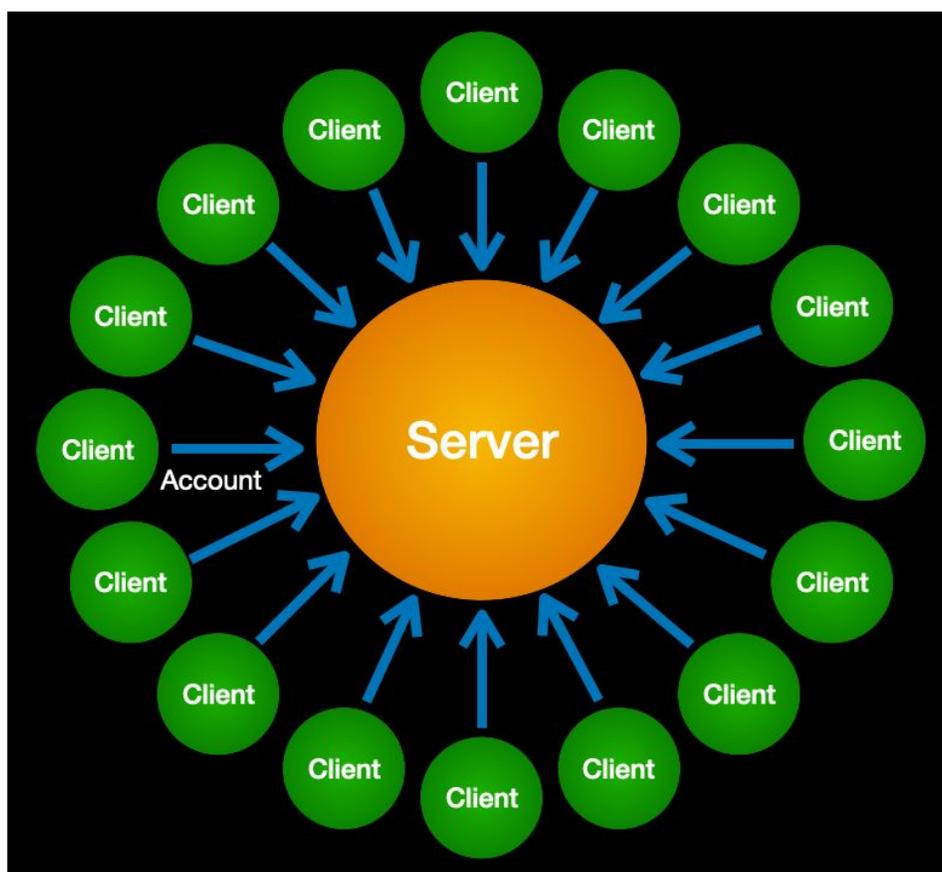
It's easy to forget that, with the Internet of today where consumers and businesses across all sectors participate in a global digital economy of unprecedented scale and complexity, *it did not start out that way*. As little as forty years ago, we were still working on individual "personal computers" that did not even connect to a network.

Then came "sneakernets", single-office networks, local-area networks (LANs), dial-up networks (CompuServe, MCI Mail, AOL), and finally the worldwide "network of networks" now known as the Internet. And at each step, the gap between how we establish trust offline and how we do it online was widening. Why?

Login Accounts: The Accidental Actor

Since the first computers were still the size of refrigerators, access to the “login” terminals worked like everything else in the real world: a door with a guard who would let you in after verifying the credentials you carry in your wallet. But as soon as we moved into a networked world, *login access could no longer be controlled via physical security*. So we created **computer-based access controls** to the login account. This was the birth of the dreaded username and password.

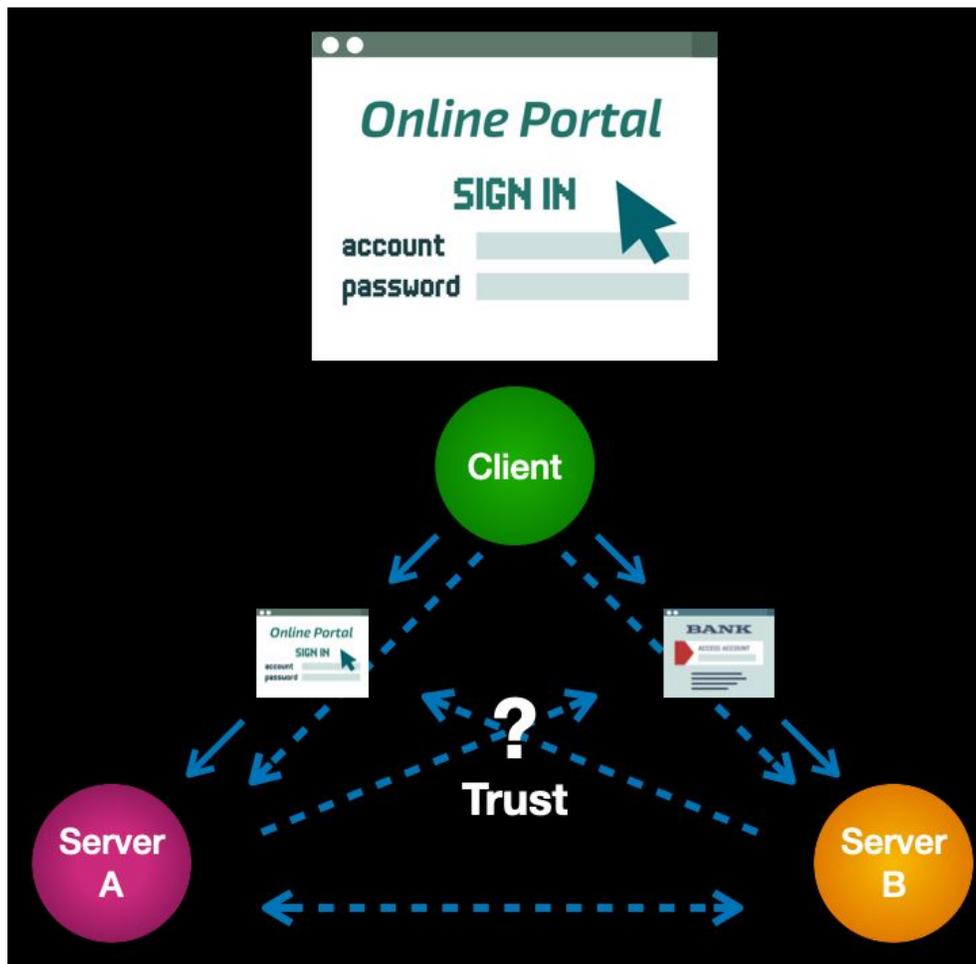
Because the login account was the virtual “door” to a server—and the server the gateway to a network—we tried to control everything by virtue of login accounts. The supremacy of servers in this “client-servitude” model is shown below.



“If the only tool you have is a hammer, everything begins to look like a nail.” Login accounts became the accidental actor in the middle of all our online interactions. No network task requiring trust could be performed without one or more logins. Pretty soon we had hundreds of usernames and passwords, and the trust gap widened even further.

Federating Accounts

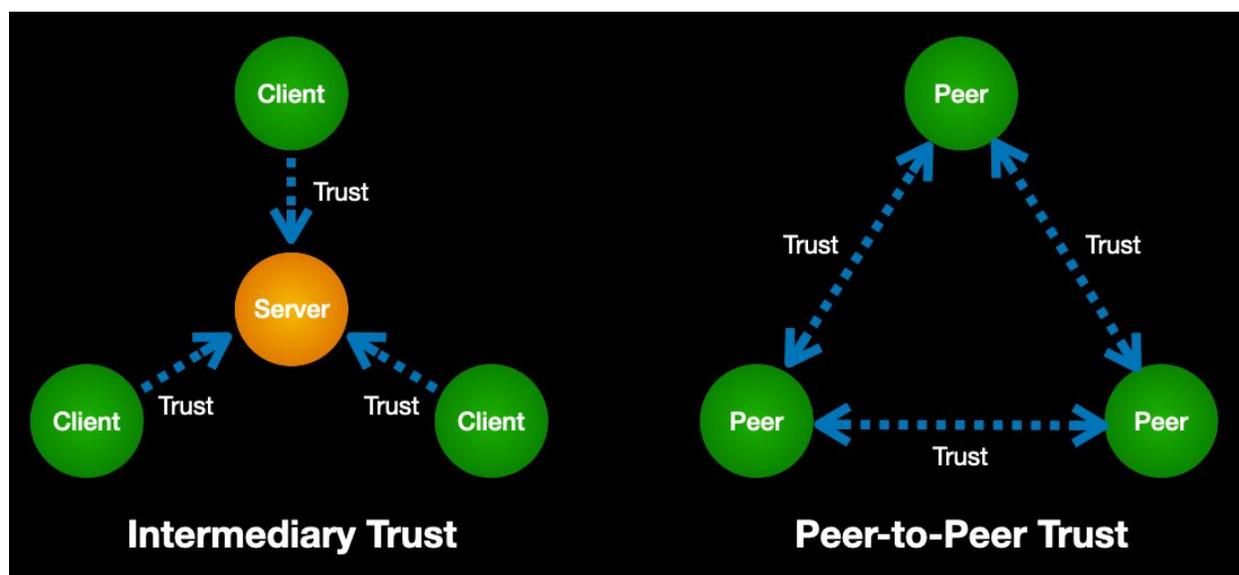
Faced with this escalating problem, technologists unwittingly drove the next wedge into the trust gap by trying to swing the server hammer even harder. They created the “federated login”, where you could reuse the account you had with one server to login to another server. This gave rise to the “single sign-on portal” that is now ubiquitous on most intranets. The mass-market equivalent is social login buttons—from Facebook, Google, LinkedIn, Twitter, etc.—that you see on many consumer-facing websites.



Federated login protocols like SAML and OpenID Connect do relieve some of the pain of maintaining long lists of usernames and passwords and repeatedly logging into sites and services. So they have achieved some measure of adoption. However there is a simple structural reason that they have not solved the ever-widening trust gap.

The Problem with Intermediaries

It is easy to spot the fundamental problem with intermediaries by looking at the **trust model**—how trust actually flows between the parties. In the current account-based client-server paradigm, all trusted interactions must be mediated by a server—and all parties must be integrated with that server. Whoever controls this server must be trusted by all the parties to the interaction. This is the model shown on the left below.



Contrast this with the peer-to-peer trust model on the right. No intermediaries needed. No server integration needed. Every peer forms trust relationships directly with every other peer. Each peer determines its own policies for trusting another peer.

This, ironically, is exactly how the trust model for real-world credentials work. Each peer is a holder of its own credentials and a verifier of another peer's credentials. Any peer can be an issuer of credentials when needed.

This is the root cause of our trust gap. Our current Internet trust model requires intermediaries that are not natural in a decentralized, peer-to-peer trust model. What's worse, the prevailing Internet "surveillance economy" business model incents these intermediaries—otherwise known as "platforms"—to monetize these private interactions, creating significant privacy issues and further widening the trust gap.

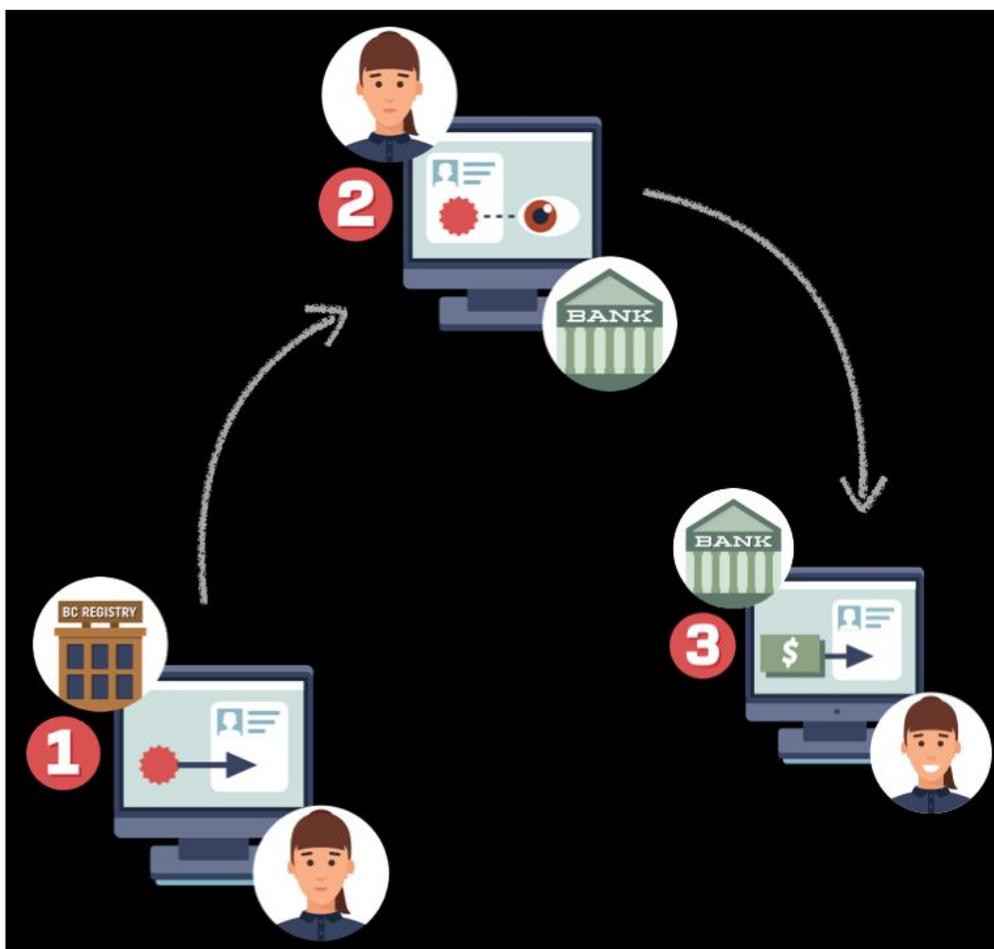
No such incentives exist for the peer-to-peer trust model. So how do we reclaim it?

Part Three: The New Era of Digital Trust

Back to the Future with Digital Credentials

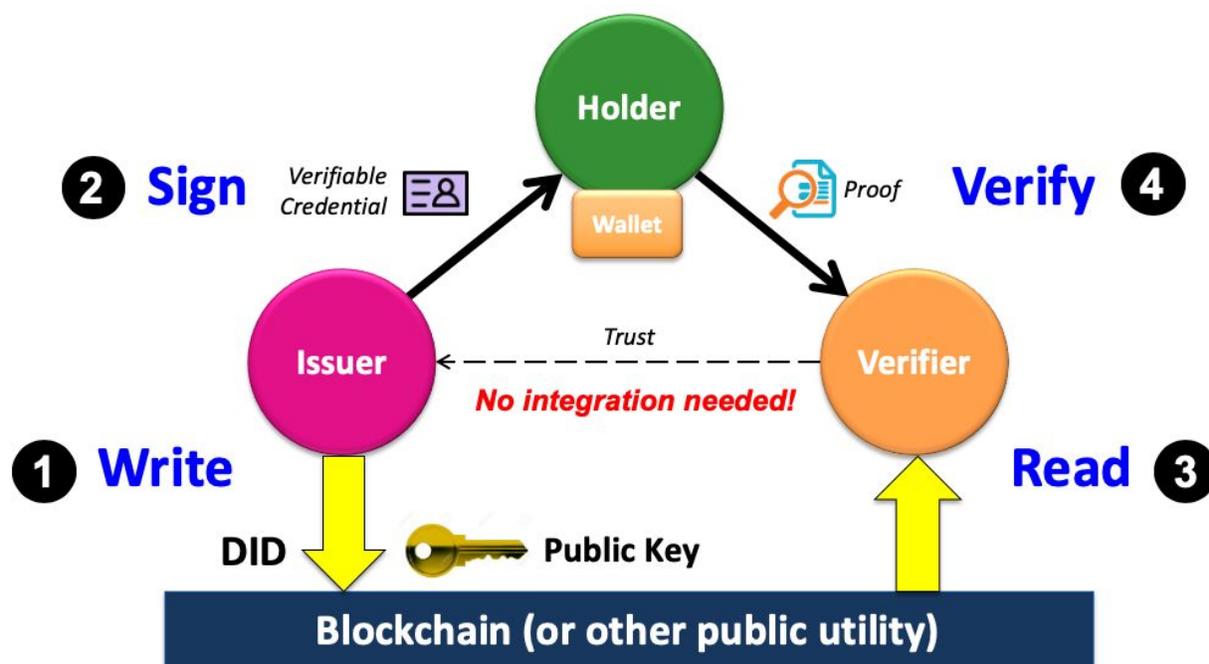
There are obvious reasons we didn't immediately port our long-established real-world trust model of physical credentials to the digital world. Physical credentials are both relatively easy to produce (via conventional printing/stamping technology), and relatively easy to verify (via human inspection, if we accept a reasonable degree of error). Digital credentials are much harder. They were a bridge too far when the Internet was young.

But now that is maturing, the benefits of introducing digital credentials would be enormous. Each of us could obtain credentials in a digital wallet just like we obtain physical credentials today. Imagine how much simpler the journey would be for a business owner like Sally, shown below. In step one she could obtain a digital license for her business. In step two she could take that credential to a bank to open a business bank account. In step three she can take both the business license and banking credential to another government agency to obtain a small-business loan—all online.



The Verifiable Credential Trust Triangle

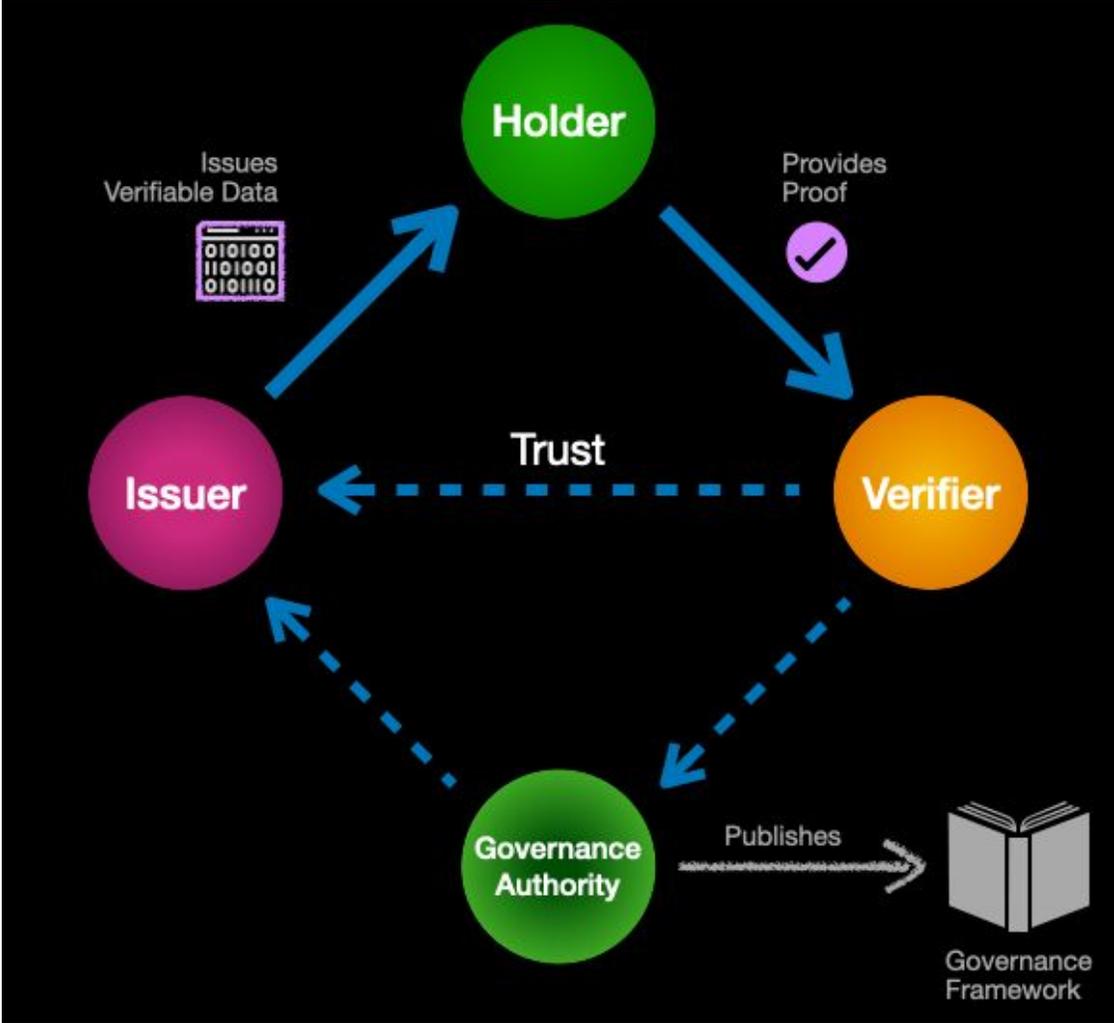
Thankfully the promise of digital credentials was recognized several years ago by pioneers at the World Wide Web Consortium (W3C). They began the effort to standardize the file formats and digital signatures needed. The result was [the Verifiable Credentials Data Model 1.0 specification](#), approved as a full W3C standard in September 2019. Below is a diagram showing how verifiable credentials work.



1. First the issuer writes a Decentralized Identifier (DID) together with its public key (and any other cryptographic material needed for the issuer's verifiable credentials) to a blockchain (or other sufficiently trusted public utility).
2. Second, the issuer uses its private key to digitally sign a verifiable credential it issues to a qualified holder, who stores it in her own digital wallet. Note that for privacy preservation, this entire issuance process takes place **off-chain**.
3. Third, a verifier requests a digital proof of one or more credentials from the holder. If the holder consents, the holder's wallet generates and returns the proofs to the verifier. Since the proofs contain the issuer's DID, the verifier uses it to read the issuer's public key and other cryptographic data from the blockchain.
4. In the final step, the verifier uses the issuer's public key to verify that the proofs are valid and that the digital credential has not been tampered with.

Economy-Scale Digital Trust

With verifiable credentials and digital wallets, we can use the same trust model—and mental model—as we use with physical credentials and wallets. Furthermore, we can use governance frameworks *to adapt this model to any trust community and scale it to any size trust network*. This **digital governance trust triangle** is shown below.

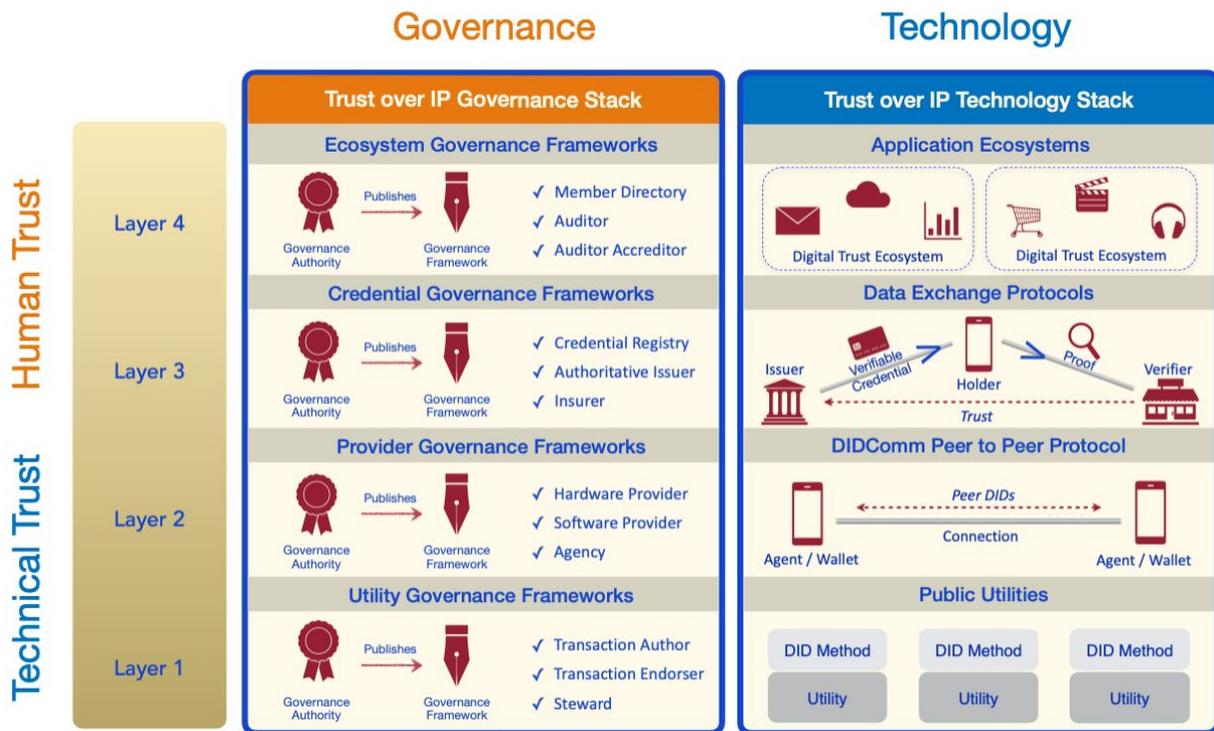


As this diagram suggests, digital governance frameworks are the backbone of this new era of digital trust. Every digital credential in your wallet should be backed by a governance framework that spells out the business, legal, and technical rules under which that credential operates. By combining the technical trust of W3C Verifiable Credentials and DIDs with the human trust codified in these governance frameworks, we can finally usher in a new era of Internet-scale digital trust infrastructure.

Part Four: The Trust over IP Stack

The Dual Stack Design

As developer communities began implementing DIDs and verifiable credentials, they recognized this new peer-to-peer trust model could underpin an entire layer of Internet-scale digital trust infrastructure. As is usually the case, their initial efforts focused primarily on proving out the technology side of the stack. But as these technical solutions started bearing fruit, customers began coming to the table looking for real-world solutions. That’s when attention turned to the “other half” of the stack—the practical governance and policy questions that must be answered in order to drive business, legal, and social acceptance. The result is the **dual stack** shown below.

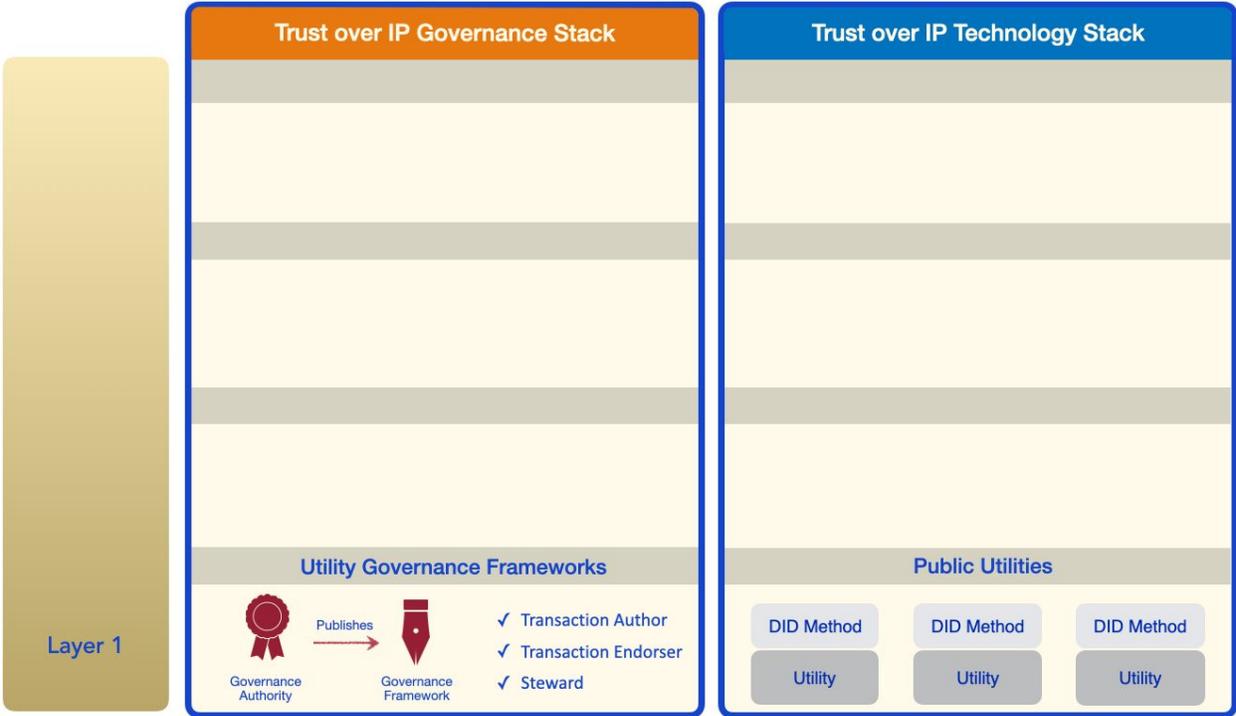


Whereas early versions of the ToIP stack reflected its historical origins—technology on the left followed by governance on the right—real-world experience soon taught us to reverse it. *Governance first*. In other words, implementing ToIP-based solutions should begin with **business requirements**, then move to **policy requirements** transparently communicated in governance frameworks. Only then should you choose the technology components required to implement those policies.

Layer One: Public Utilities

The first two layers of the ToIP stack are designed to provide **technical trust**—the assurance that one machine can establish a secure, private connection with another machine. To do this using [public key cryptography](#), you must be able to strongly verify the **public key** of the party you are connecting to. The [W3C Decentralized Identifier \(DID\) specification](#) solves this problem without using centralized [certificate authorities](#) by standardizing how you can permanently identify and verify a public key stored on a blockchain or other distributed system.

This solution gives rise to public utilities that serve as strong **cryptographic roots-of-trust** for the DIDs and public keys of verifiable credential issuers. ToIP Layer One utilities can be implemented using any technology that can provide the necessary trust assurances, e.g., blockchains (of any kind), distributed ledgers, decentralized file systems, distributed hash tables, and so on.

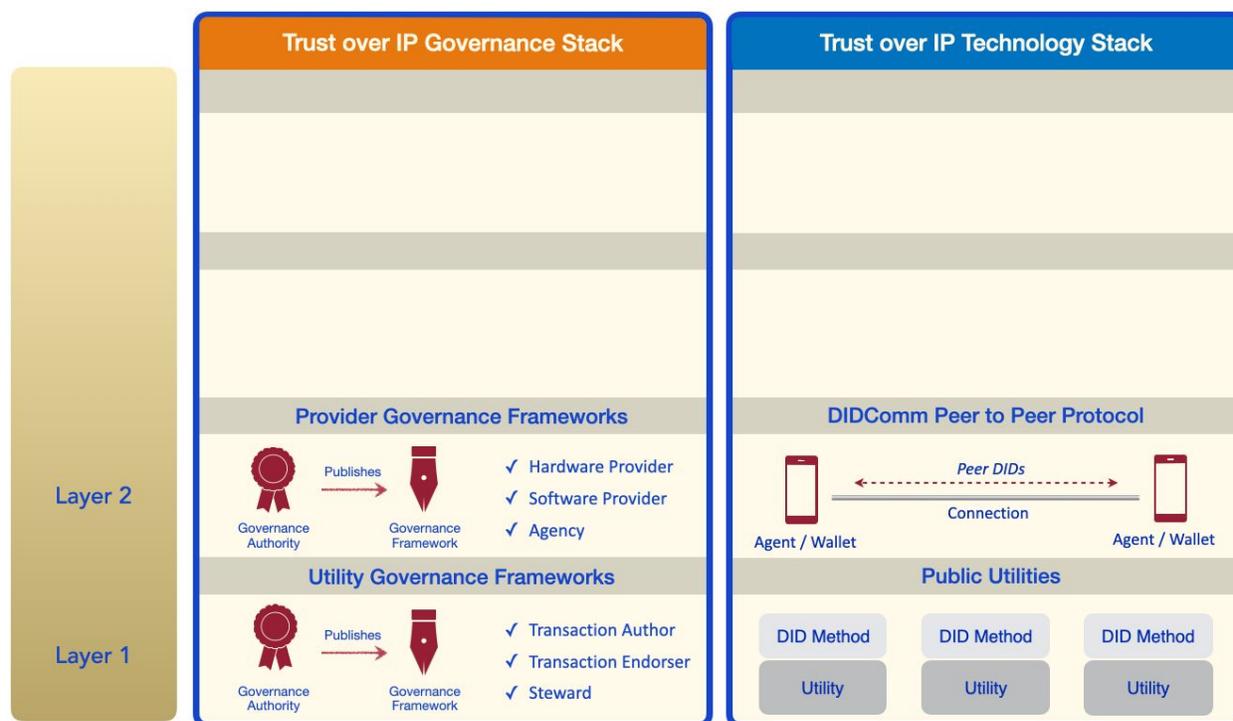


Although technical trust is machine-to-machine, implementing technical trust still requires humans to design, code, test, and certify these systems. This is the job of Layer One **utility governance frameworks** that specify the policies under which a utility is implemented and operated such that it can be trusted by the higher layers.

Layer Two: DIDComm Peer-to-Peer Protocol

If Layer One is about the strong cryptographic *roots* of technical trust, then Layer Two is about the *branches*—the digital wallets and digital agents needed to form secure, private peer-to-peer connections using either public DIDs (from Layer One) or [peer DIDs](#). The latter are exchanged directly between the peers and never need to touch a blockchain—a significant advantage for both scalability and privacy.

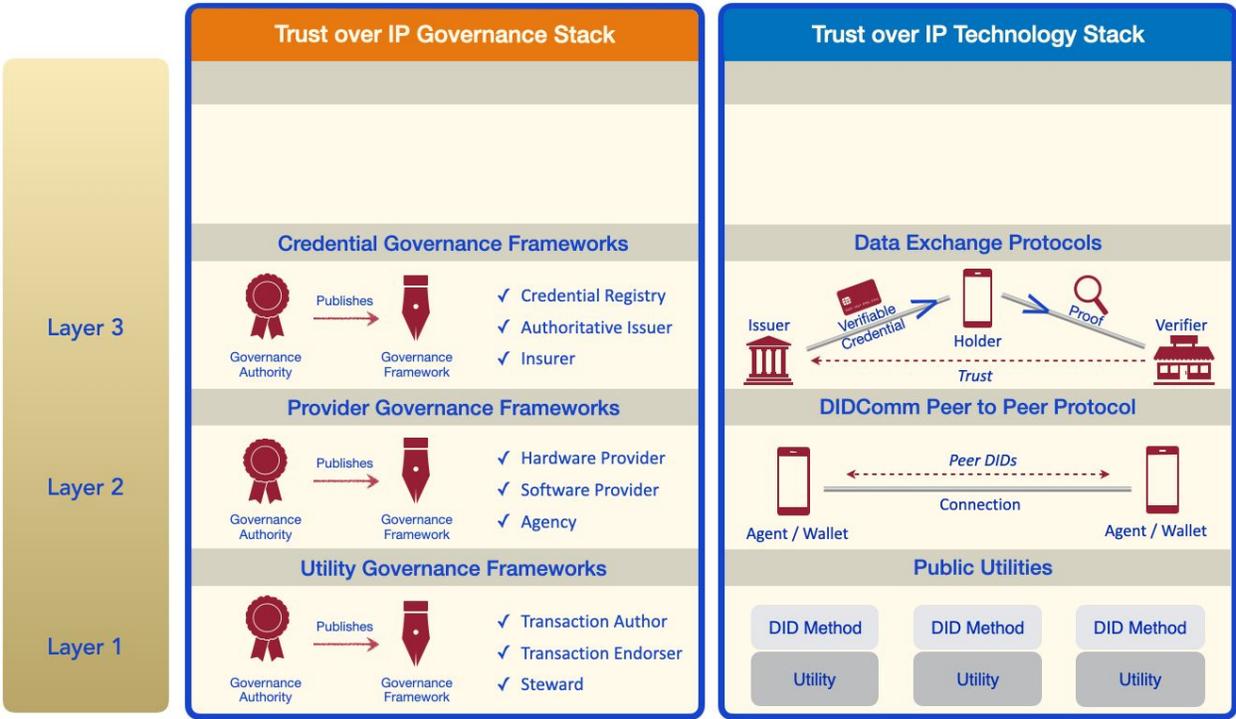
Just as the Internet Protocol (IP) forms the [narrow waist](#) of the TCP/IP stack that powers the Internet, the DIDComm protocol (currently a Working Group at the [Decentralized Identity Foundation](#)) forms the narrow waist of the ToIP stack.



Again, although technical trust is machine-to-machine, how digital wallets and agents are actually implemented makes a tremendous difference not only to the security and privacy of users, but to their confidence that their personal data and credentials are truly portable and vendor-independent (unlike the proprietary digital wallets built into our smartphones today). This is the province of **provider governance frameworks** that can specify the privacy, security, and data protection standards against which hardware providers, software providers, and cloud hosting providers can be certified.

Layer Three: Data Exchange Protocols

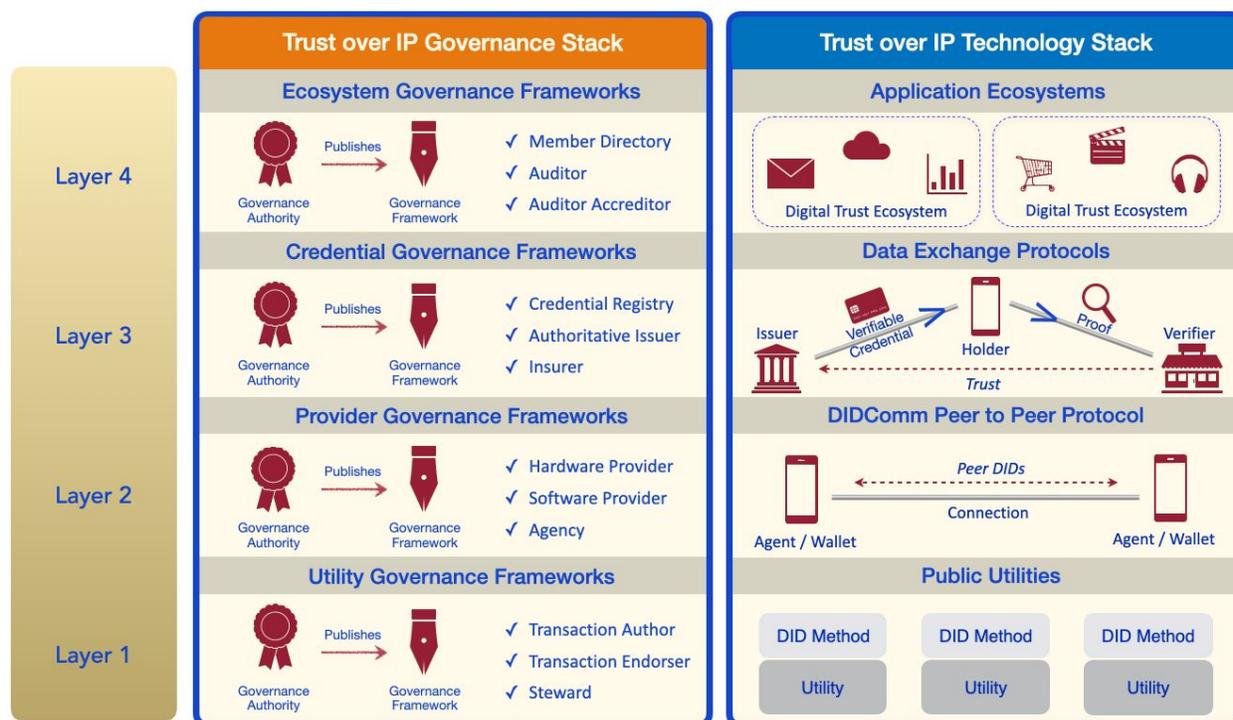
Layers Three and Four are where **human trust** is established and maintained. On the technical side of the stack, Layer Three is the home of the **verifiable credential trust triangle** discussed in Part Three. This is the layer where issuers, holders, and verifiers exchange credentials and proofs using **credential exchange protocols** that run on top of DIDComm. Note that these are just one example of the kind of trusted data exchange protocols that can operate at Layer Three—many other types of secure messaging and workflow automation protocols can be implemented at this layer.



On the governance half of the stack, Layer Three is where the **governance trust triangle** comes into full play. Almost any digital credential that will be issued by multiple issuers and/or accepted by a wide range of verifiers needs a **credential governance framework**. It will define what issuers will issue what credentials under what policies to what holders with what level(s) of assurance—and under what trust mark(s). This is the information verifiers need to make their own trust decisions about relying on a proof from the credential—just as the Mastercard operating rules tell merchants exactly what they can expect when accepting a Mastercard.

Layer Four: Application Ecosystems

Layer Four is the application layer—the layer where humans interact with applications in order to engage in trusted interactions that serve a specific business, legal, or social purpose. Just as Internet-enabled applications call the TCP/IP stack to communicate over the Internet, ToIP-enabled applications call the ToIP stack to register DIDs, form connections, obtain and exchange verifiable credentials, and engage in trusted data exchange using the protocols in Layers One, Two, and Three.



Layer Four is specifically designed to enable **digital trust ecosystems**—entire families of applications and credentials that are not only designed to interoperate technically, but which share a common **ecosystem governance framework**. This specifies the purpose, principles, and policies that apply to all governance authorities and governance frameworks operating within that ecosystem—at all four layers of the ToIP stack.

An ecosystem governance framework can enable nearly frictionless data exchange between apps, sites, and businesses while providing a consistent user experience of security, privacy, and data protection across the ecosystem that can be as important to consumer confidence as a consistent user experience of the controls for driving a car (steering wheel, gas pedal, brakes, turn signals) are to driver safety around the world.

Part Five: The Role of the Trust over IP Foundation

The Linux Foundation

The founders of the Trust over IP Foundation chose the Linux Foundation (LF) as our home for the simple reason that it hosts the largest and most successful open source projects in the world. Besides Linux itself, the LF hosts over 240 independent projects including the Cloud Native Computing Foundation, Automotive Grade Linux, Carrier Grade Linux, the R Consortium, the Node.js Foundation, and the GraphQL Foundation.



16B USD

Estimated development cost of the 100+ world's leading projects hosted at The Linux Foundation



35,000

Technologists attend our events annually, from more than 11,000 companies and 113 countries



1 Million

Open source professionals have enrolled in our free open source training courses



10 / 10

Largest cloud service providers are Linux Foundation project contributors and members

Secondly, the LF is already the home of two directly related peer projects:

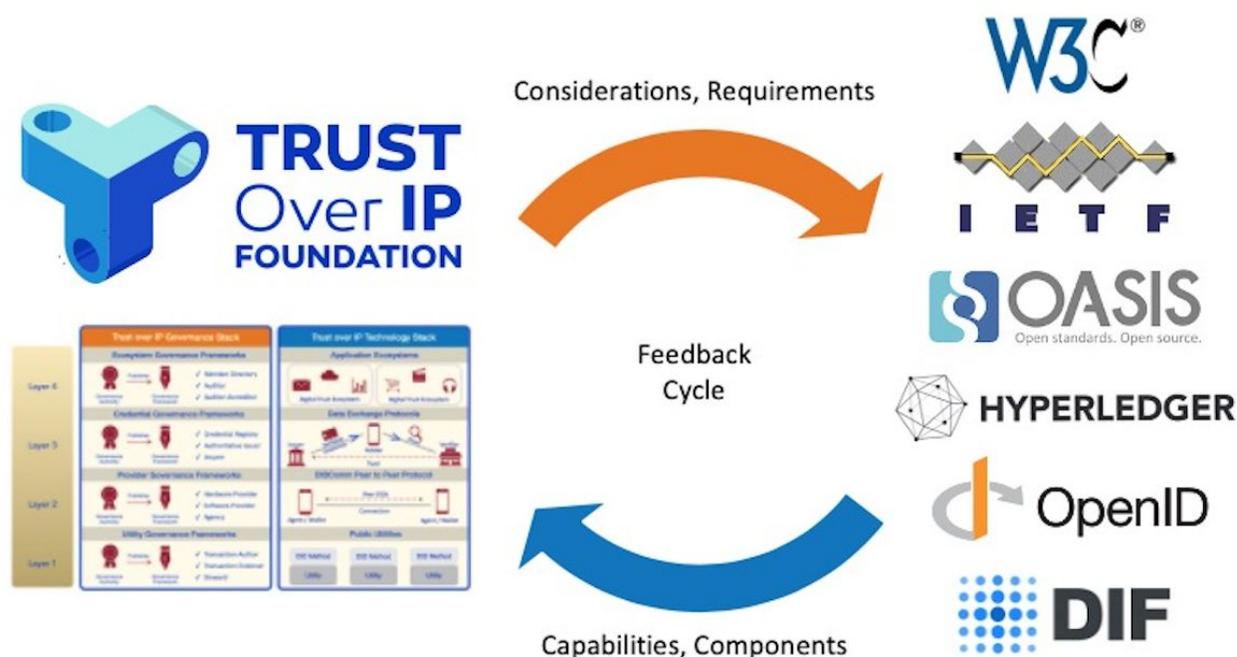
1. [Hyperledger](#), the umbrella organization hosting over a dozen projects for advancing blockchain technology for business. Three Hyperledger projects—[Indy](#), [Ursa](#), and [Aries](#)—implement key components of the ToIP stack.
2. [Decentralized Identity Foundation](#) (DIF), a membership organization building foundational components of open, standards-based decentralized identity. DIF is currently the home of the [DIDComm Working Group](#)—the “narrow waist” protocol at the heart of Layer Two of the ToIP stack. DIF also hosts several other Working Groups focused on DIDs, secure data storage, and other aspects of the stack.

Mission of the Trust over IP Foundation

The charter of the Trust over IP Foundation is to:

Define a complete architecture for Internet-scale digital trust that combines cryptographic trust at the machine layer with human trust at the business, legal, and social layers.

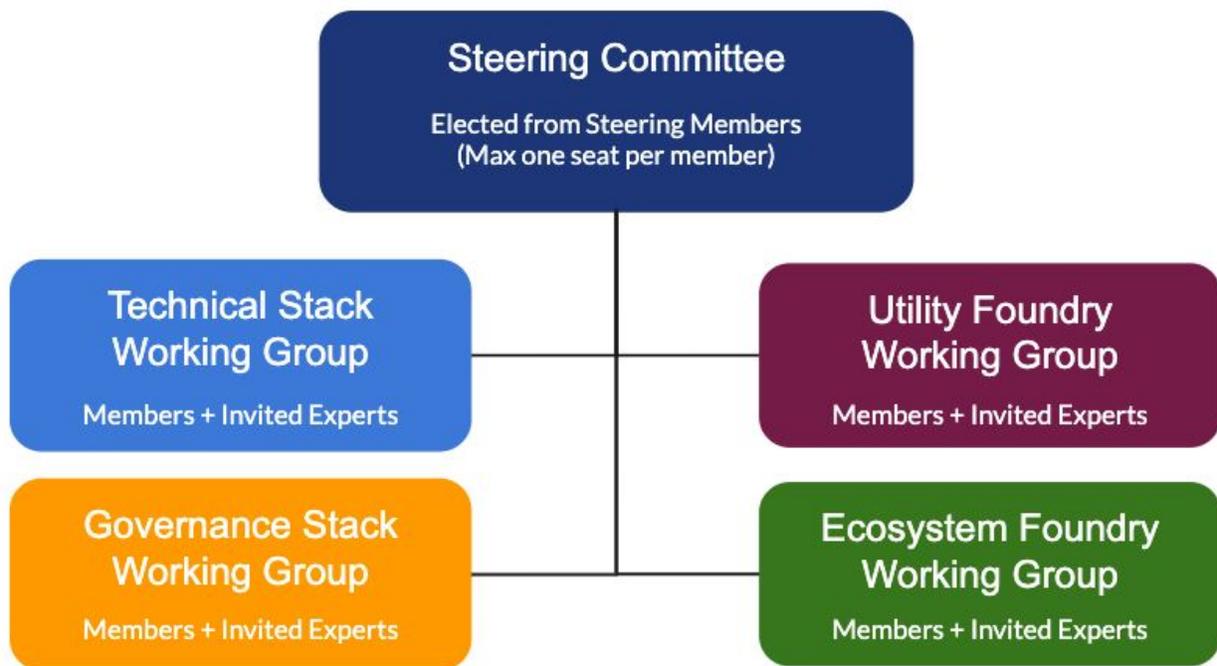
Note that this mission is not to develop all of the standards or components included in the ToIP stack—rather it is to specify how these elements can be combined to fulfill the requirements of all four layers of the stack, for both governance and technology. This means the ToIP Foundation will work closely with other standards development organizations (SDOs), industry foundations, and other consortia to combine their open standards, architectures, and protocols into a complete and coherent stack for Internet-scale digital trust infrastructure.



Note that the organizations listed are those whose work is either referenced by or contributing to some portion of the ToIP stack. We expect this roster of relationships to grow as our work on the ToIP stack advances through the ToIP Foundation Working Groups.

Governance and Working Groups

Like almost all Linux Foundation projects, the Trust over IP Foundation is governed by a Steering Committee composed of representatives of the Steering Members. Steering Membership is available at two levels: >100 employees (USD \$20K/yr), and <100 employees (\$5K/yr). Associate Membership is also available at half that cost, and Contributor Membership are available to both individuals and organizations at no cost.



The work of the Foundation will proceed in four initial Working Groups:

1. The **Technical Stack Working Group** will define the specifications and interoperability testing requirements for the ToIP Technology Stack.
2. The **Governance Stack Working Group** will define the models, templates, guidelines, and recommended best practices for the ToIP Governance Stack.
3. The **Utility Foundry Working Group** is a community of practice for governance authorities implementing ToIP Layer One public utilities—whether as LF projects or as external governance organizations in any jurisdiction.
4. The **Ecosystem Foundry Working Group** is a community of practice for governance authorities seeking guidance and support in implementing ToIP Layer Four digital trust ecosystems.

Join Us

As of our formal launch on 5 May 2020, the Trust over IP Foundation has 29 Founding Members—17 Steering Member and 12 Contributor Member organizations. Please visit our website at <https://trustoverip.org/> to see the full membership list. New members are welcome at any time—whether the Steering or Associate paid levels, or at the Contributor level at no cost.

It is very important that there be no barrier of entry to any individual or organization who wishes to contribute to the development of the ToIP stack.

Please contact us if you have any questions.

Together we look forward to building the trust layer for the Internet.



TRUST
Over IP
FOUNDATION

© 2020 Trust over IP Foundation. This work is licensed under Creative Commons AttributionShareAlike 4.0 International License (<http://creativecommons.org/licenses/by-sa/4.0/>).