TRUST Over IP FOUNDATION

Introduction to Trust Over IP

Version 2.0

17 November 2021

This publicly available whitepaper was approved by the ToIP Foundation Steering Committee on 17 November 2021.

The mission of the Trust over IP (ToIP) Foundation is to define a complete architecture for Internet-scale digital trust that combines cryptographic assurance at the machine layer with human accountability at the business, legal, and social layers. Founded in May 2020 as a non-profit hosted by the Linux Foundation, the ToIP Foundation has over 300 organizational and 100 individual members from around the world.

Please see the end page for licensing information and how to get involved with the Trust Over IP Foundation.

Table of Contents

Document Information

Contributors

Carly Huitema Daniel Bachenheimer — Accenture Darrell O'Donnell — Continuum Loop Drummond Reed — Evernym Jacques Bikoundou Judith Fleenor — Trust Over IP Foundation Kaliya Young — COVID-19 Credential Initiative Karen Hand — Precision Strategic Solutions Karl Kneiss — IdRamp	John Jordan — Province of British Columbia Lynn Bendixsen — Indicio P. A. Subrahmanyam — CyberKnowledge Sankarshan Mukhopadhyay — Dhiway Networks Scott Perry — Scott S. Perry CPA, PLLC Victor Syntez Vikas Malhotra — WOPLLI Technologies Wenjing Chu — Futurewei

Acknowledgements

We would also like to thank Peter Stoyko of Elanica for his design of the ToIP logo and his wonderful interactive graphics of the ToIP stack.

Revision History

Version	Date Approved	Revisions
1.0	10 May 2020	Initial Publication
2.0	17 November 2021	Second Edition

Terms of Use

Introduction

On 5 May 2020, the Linux Foundation announced a new addition to its roster of global open source/open standard/open governance projects: the **Trust Over IP Foundation** (ToIP).

The mission of this Foundation is to simplify and standardize how trust[1] is established over a digital network or using digital tools (whether online or disconnected). The goal is to create a safe and private space for all digital interactions—whether between individuals, businesses, governments, or any type of "thing" we might digitally interact with.[2] The primary tool for achieving this objective is a "stack" roughly analogous to the TCP/IP stack that powers the Internet. However, this stack is *not just technology*, rather it combines cryptographic verifiability at the machine layers with human accountability at the legal, business, and social layers.

This approach has resonated strongly with many stakeholders in the digital economy. In its first year, the ToIP Foundation grew from 27 to over 300 member organizations, doubled the number of its Working Groups, and been invited to host or contribute to a number of national and international projects focused on building trust online, including the Good Health Pass Collaborative, the UK Digital Identity Trust Framework, the COVID-19 Credentials Initiative, and the Ontario Digital Market Consultation.

In this whitepaper, we will first provide a summary of ToIP. Then we will dive into the question of why trust between parties intermediated by digital technology has been such a hard problem to solve, and why centralized and federated identity systems and public key infrastructures have not been able to overcome these challenges.

We then turn to the proposed solution by analyzing the trust model that underlies the physical credentials we use in the real world today. We will show how, by following a few fundamental design principles, this model can be adapted to the digital world. We will explain how this approach evolved into the four-layer, dual-sided architecture of the ToIP stack.

Next, we will introduce the ToIP Foundation as a collaborative community that drives the design, development, and adoption of digital systems that follow the dual-sided ToIP stack architecture. We will cover the structure and governance of the Foundation, its current Working Groups, and how to get involved.

We will conclude with a look at our roadmap going forward to see where we, as a community, are aiming to be in one, three-, and five-years' time.

---

[1] The term "trust" has been the subject of entire papers and books. This subject is currently being explored by the ToIP Foundation on two fronts: 1) an analysis of digital trust relationships in the context of the ToIP stack in a deliverable called Design Principles for the ToIP Stack, and 2) the ToIP Concepts and Terminology Working Group is defining terms and mental models associated with digital trust.

[2] In the context of ToIP, a digital "thing" is much broader than the Internet of Things (IoT). IoT deals with devices of any kind connected to the Internet so they can send and receive data. The universe of digital things in our lives is much broader and includes all the digital documents, files, photos, videos, software programs, bots, and even AI programs that we interact with, plus digital twins of physical objects.

ToIP in a Nutshell

At a high level, the goal of the ToIP Foundation can be expressed in one sentence:

> We develop tools and specifications to help communities of any size use digital networks to build and strengthen trust between participants.

"Trust" can be difficult to define in the abstract but is quintessentially a human belief that is always *relational* (between two parties), *directional* (going in one direction or the other), and *contextual* (applying in a specific context).[3] So the "atomic unit" of trust is a single relationship in a single direction in a single context as shown in Figure 1. Examples include a parent trusting a student to be a babysitter, or a company trusting a contractor to fix a particular type of machine.

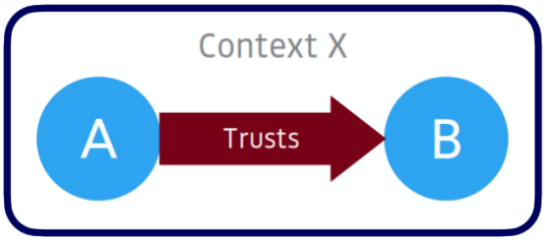

Figure 1. A direct trust relationship in one direction in one context

Trust is also an inherent component of any kind of socio-cultural ecosystem that connects people to other people within a community. Countries, cities, companies, churches, partnerships, schools, social networks—all are communities of some kind, and in order to function successfully they all depend on various degrees of trust between participants.

As these various social structures grow, they run into the problem of how to *scale trust*. For example, as a community becomes more geographically dispersed, it becomes increasingly difficult for every member to build a direct trust relationship with every other member they wish to interact with. So, we need ways to enable *transitive trust*—turning trust between two members into trust with a third.

The basic structure of transitive trust can be viewed as the *trust triangle* shown in Figure 2. On the left, A has a direct trust relationship with B, and B has a direct trust relationship with C. If both of those are in the same context X, A can then have some degree of transitive trust in C as shown on the right.[4]

[3] For an in-depth examination of the seven core principles of trust, please see Design Principles for the ToIP Stacks, one of the first deliverables from the ToIP Foundation.

[4] Transitive trust is neither automatic nor absolute. In other words, the right half of Figure 2, A *may* decide to trust C to some *limit*. It does not mean A automatically trusts C to the same limit that B does.

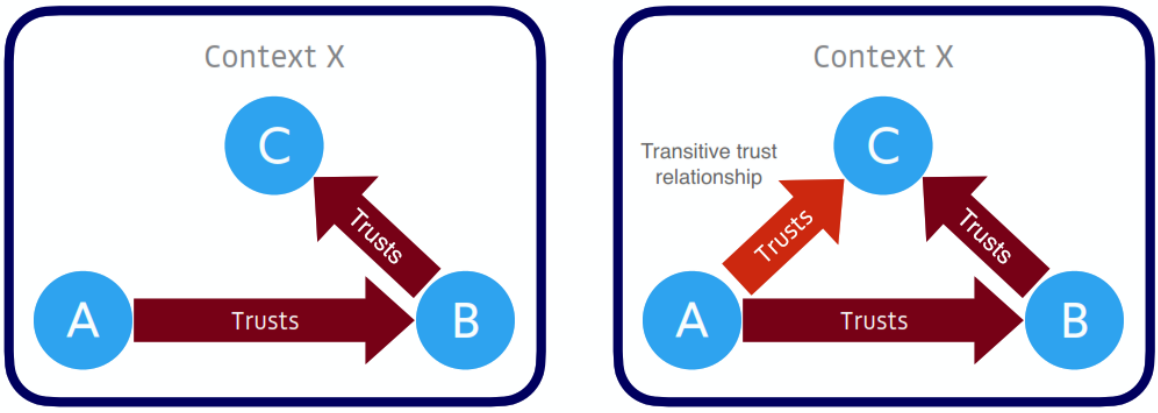

Figure 2. Two connected trust relationships in the same context (left), can lead to a transitive trust relationship completing the "trust triangle" (right)

For example, if you (A) trust your friend Bob (B) to refer a babysitter and Bob recommends Cathy (C), then you are likely to trust Cathy to be a good babysitter. Some of your trust in Bob has been transferred to Cathy by referral rather than gained through direct experience. The same is true in an organizational context, e.g., if Acme (A) trusts Bonfix (B) to fix copy machines, and Bonfix (B) trusts CopyCity (C) with regard to office machines, then Acme is likely to trust CopyCity to fix its copy machines—but maybe not to fix its delivery trucks.

In society today we have hundreds of ways to convey transitive trust. Governments, currencies, trademarks, diplomas, licenses—all of these are instruments that help us leverage the trust decisions made by other people or organizations. In fact, every credential you carry in your wallet is a tool for transferring some of the trust the *issuer* has in you, the *holder*, to a third party, the *verifier*. It helps the verifier to make some type of trust decision about you—to board a plane, rent a car, enter a building, make a purchase, etc.

As our economies, societies, and governments become increasingly reliant on digital intermediation, the question is how can we enable transitive trust "at a distance", i.e., over *digital networks*? As this white paper will explain, so far, we have been trying to do this primarily by leveraging the digital accounts you have with large service providers—for example, by using social login buttons to log into websites. But this requires using the same intermediaries in all our online trust relationships. As a mechanism for transitive trust, this approach is unnatural, inefficient, and privacy-invasive. By inserting these digital intermediaries, individuals are exchanging convenience for privacy, exposing their personal information to organizations in ways that put them at risk—risk of surveillance, risk of their data being shared or sold with other organizations, risk of data breaches. As our lives become increasingly digital, these risks become both broader (more incidents) and more acute (greater damage).

What we need are new mechanisms that enable transitive trust to be conveyed quickly and easily, between any set of peers, in any context—the way it works in the real world. With these mechanisms, trust could evolve more quickly and organically than by using hierarchical, "top-down", centrally-controlled systems. It can give rise to what we call *digital trust ecosystems*. Just like a spider's web consists of thousands of individual strands, a digital trust ecosystem can consist of thousands or millions of trust triangles between the parties who are incented to build trust relationships to accomplish their mutual goals. Figure 3 is a visualization of an ecosystem for digital health "passes" where the individual, at the center, has a digital credential as the result of a

vaccination (syringe) or a test (lab machine) which can be used to prove their health status to a healthcare provider (ambulance), entertainment venue (arena), airline (airplane), or a school (blackboard).[5]

Figure 3. Interconnected stakeholders in a health pass digital trust ecosystem

Just as the real world consists of millions of interconnected ecosystems of all types and sizes, so will the world of digital trust ecosystems. The goal of ToIP is to make digital trust ecosystems accessible and interoperable across communities of any scale:

- International-scale ecosystems like the Good Health Pass for travellers.
- Country-wide ecosystems serving governments and their citizens.
- Industry-wide ecosystems supporting supply chains for goods and services.
- Regional or city-wide ecosystems serving specific populations or markets.
- Local ecosystems serving small businesses and communities, e.g., gym memberships, museums, co-ops, churches, farmer's markets.

Most importantly, just like real-world living ecosystems, digital trust ecosystems can emerge, grow, and connect organically—*they do not require permission from central authorities or gatekeepers*. They can be independently organized at any scale.

[5] This figure is from the interactive graphic of the ToIP stack on the ToIP website—a recommended tool for exploring ToIP. Health passes were the core focus of the Good Health Pass Interoperability Blueprint, a ToIP recommendation for COVID-19 health credentials, which was published in August 2021. Note that the Good Health Pass recommendations apply *only* to international travel and not to the other use cases depicted in this graphic.

Our digital trust landscape today is one where digital signatures, encryption, and public key infrastructures (PKIs) are predominately centrally-managed, top-down structures that can be expensive to set up and maintain and require significant technical expertise to operate. The mission of the ToIP Foundation is to not only turn this situation around, but to help expand the opportunities of digital trust by bringing **a unified stack of standards for technical interoperability**—the same approach that has been successful with the Internet and the Web—together with **a unified model for expressing the rules and policies** ("governance") by which people and organizations can cooperate to achieve trust.

Why Has Digital Trust Become Such a Major Problem?

In his ground-breaking series of essays published in 2004 called *The Laws of Identity*, Kim Cameron—Microsoft's Chief Architect for Identity from 2004 to 2019—said:

> *The Internet was built without a way to know who and what you are connecting to. This limits what we can do with it and exposes us to growing dangers. If we do nothing, we will face rapidly proliferating episodes of theft and deception that will cumulatively erode public trust in the Internet.*

Kim's prophecy has come frighteningly true despite 20 years of collective work trying to solve these problems. The onslaught of reports and statistics about the breakdowns in digital trust has become mind-numbing. Following is just a sampling.

Digital Trust Statistics

Passwords

- Regular Internet users have an average of 85 passwords for all their accounts. (Cnet, 2020)
- The most used password in the world remains **123456** followed by **123456789**, **qwerty**, **password**, and **12345**. (Cybernews, 2021)
- 80% of all hacking incidents are caused by stolen and reused login information. (Verizon, 2020)

Phishing

- As of 2020, phishing is by far the most common attack performed by cyber-criminals, with the U.S. FBI's Internet Crime Complaint Centre recording over twice as many incidents of phishing than any other type of computer crime. (FBI Internet Crime Complaint Centre, 2021)
- Google has registered 2,145,013 phishing sites as of Jan 17, 2021. This is up from 1,690,000 on Jan 19, 2020 (up 27% over 12 months). (Tessian, 2021)

Data Breaches

- There were 1,767 publicly reported data breaches in the first six months of 2021, which exposed a total of 18.8 billion records. (Risk Based Security, 2021)
- Over 90% of all healthcare organizations reported at least one security breach in the last three years. 61% acknowledged they don't have effective mechanisms to maintain proper cybersecurity. (Frost Radar, 2020)
- In 2020 the average cost of a corporate data breach was $3.86 million. (Dice.com, 2020)

Privacy Erosion and Surveillance Capitalism

- 82% of web traffic contains Google third-party scripts and almost half of them are tracking users. (WhoTracks.Me, 2019)
- 74% of Internet users feel they have no control over the personal information collected on them. (Ponemon Institute, 2020)
- 72% of Americans report feeling that all, almost all, or most of what they do online or while using their cellphone is being tracked by advertisers, technology firms or other companies. (Pew Research Center, 2019)

Misinformation and Unverified Sources

- In 2020, only 29% of US adults said they mostly trust the news media. (Statista, 2020)
- In Q3 of 2020, there were 1.8 billion fake news engagements on Facebook. (German Marshall Fund, 2020)
- 56% of Facebook users can't recognize fake news when it aligns with their beliefs. (SSRN, 2018)

Artificial Intelligence (AI) Dangers

- 62% of the companies adopting AI are extremely concerned that it will increase their cybersecurity vulnerabilities; 57% are concerned about the consequences of their AI systems using personal data without consent. (Deloitte, State of AI in the Enterprise, 2020)
- 93% of automation technologists feel unprepared or only partially prepared to tackle the challenges associated with smart machine technologies. (Forrester, 2016)
- Only 36% of AI adopters are establishing policies or a board to guide AI ethics. (Harvard Kennedy School, 2019)
- The EU has drafted an Artificial Intelligence Act (AIA) specifically addressing transparency, privacy and security in the use of AI.[6]
- The National Institute of Standards and Technology (NIST) is beginning development of an AI Risk Management Framework (RMF)[7] to guide AI adoption for US federal agencies (where none currently exists).

[6] https://law.stanford.edu/publications/eu-artificial-intelligence-act-the-european-approach-to-ai/

[7] https://www.nist.gov/itl/ai-risk-management-framework

The Root Cause

In a five-part series of articles published in 2015 by the Washington Post of articles entitled Net of Insecurity: A Flaw in the Design, David D. Clark, an MIT scientist overseeing a meeting of engineers dealing with the first Internet worm attack, is quoted:

> *"It's not that we didn't think about security," Clark recalled. "We knew that there were untrustworthy people out there, and we thought we could exclude them."*
>
> *How wrong they were. What began as an online community for a few dozen researchers now is accessible to an estimated 3 billion people. That's roughly the population of the entire planet in the early 1960s, when talk began of building a revolutionary new computer network.*
>
> *Those who helped design this network over subsequent decades focused on the technical challenges of moving information quickly and reliably. When they thought about security, they foresaw the need to protect the network against potential intruders or military threats, but they didn't anticipate that the Internet's own users would someday use the network to attack one another.*

In other words, whether intentional or not, trust wasn't built into the original architecture of the Internet because the early designers trusted each other. Most of them were academics, computer scientists, or researchers who knew each other. Furthermore, they all needed access to expensive machines and sophisticated technical skills to even participate. So even though the Internet was designed to be decentralized without single points of failure, in its early days it was effectively a small club.

That all changed, as the Washington Post article explains, when the Internet turned out to be exponentially more successful than anyone ever imagined. And unfortunately, that meant the security, privacy, and trust problems grew exponentially right along with it.

How We Tried to Fix the Problem: Account-Based Digital Trust

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**. 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-server" model is shown in Figure 4.

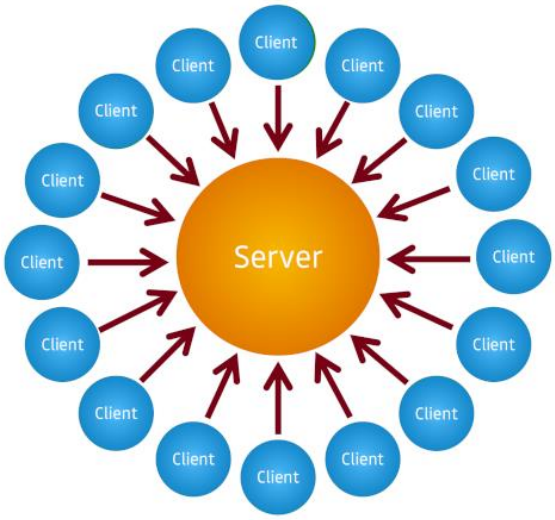

Figure 4. Servers are at the center of the client-server model of computing

"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 we began opening up a "trust gap" between the digital world and the real world—meaning a growing difference between what we do to achieve trust in our online interactions vs. how we are used to doing it in our everyday offline lives.

Federated 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"—new protocols that let you reuse the account you have 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 the social login buttons—from Facebook, Google, LinkedIn, Twitter, etc.—that you see on many consumer-facing websites.

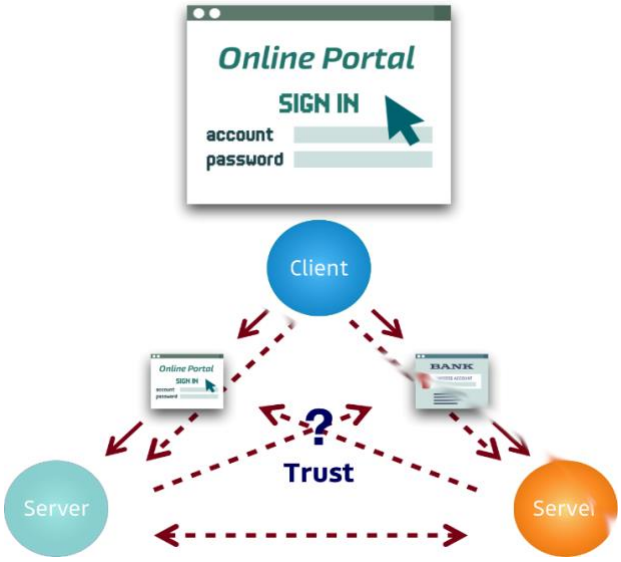

Figure 5. Using servers to mediate trust between domains is fraught with challenges

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 gap between trust in the real world and trust in the digital world.

The Fundamental 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 in Figure 6.

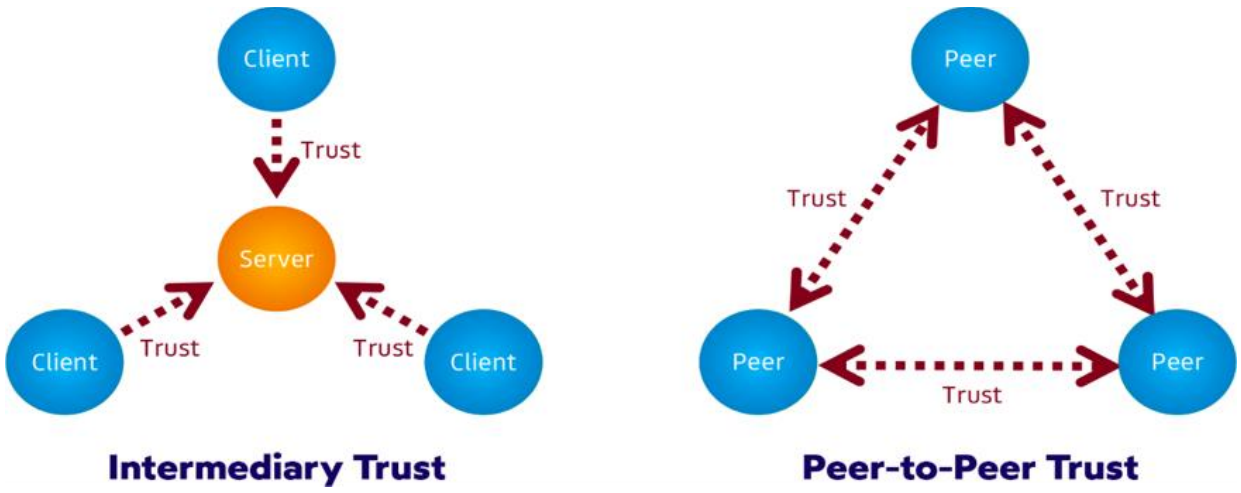

Figure 6. The intermediary trust model vs. the peer-to-peer trust model

Contrast this with the peer-to-peer trust model on the right. No intermediaries are needed. No server integration is 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 works. Each peer is a holder of its own credentials. Every peer can be a verifier of another peer's credentials. Any peer can be an issuer of credentials when needed.

In short, the peer-to-peer model is more natural, intuitive, and decentralized than the server-based model—and, if implemented with the proper cryptography and protocols, also more secure and privacy-preserving. Adopting this model would eliminate the "trust gap" that currently exists between the physical and digital worlds.

But transitioning to a peer-to-peer trust model will have its challenges for several reasons:

1. **First, federated models are well-entrenched.** 15 years of work have gone into client-server federated identity protocols such as OpenID Connect, and these have been widely deployed in some industries where they have been a definite improvement. So the transition to a peer-to-peer model will almost certainly involve hybrid approaches for some time.
2. **The prevailing advertising-based business model of the Web** provides incentives for these intermediaries—otherwise known as "platforms"—to monetize these private interactions. The peer-to-peer trust model will need incentives that do not reproduce these same privacy risks.
3. **A peer-to-peer trust model requires both standards and utilities** for decentralized digital trust infrastructure that are only just now starting to appear.

So, what will the transition to a peer-to-peer trust model look like?

The ToIP Model for Digital Trust

How Trust Works in the Real World

In-Person 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 transitive trust. We used (and still use) **credentials** of all kinds.

Figure 7. Credentials are a centuries-old solution for transitive trust

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 ensures the establishment of trust. For example, this includes:

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

These records of identity attributes are only valuable if the verifier (see Figure 8) trusts them. With physically printed documents, trust is literally in the eye of the beholder—it is conveyed via visible security features such as wax seals, colored threads, intaglio printing, color shifting inks, or holographic stamps.

The Credential Trust Triangle

The reason credentials have evolved as a universal mechanism for establishing real-world trust is the fundamental "trust triangle" illustrated in Figure 8. Note that this is an example of exactly how transitive trust works as we covered in Figure 2.

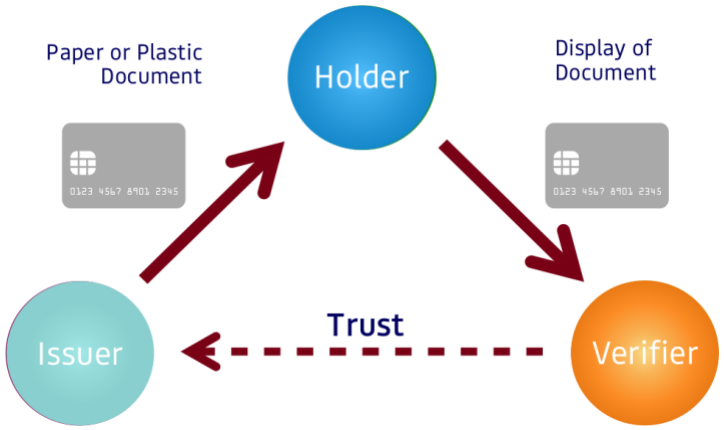

Figure 8: The three roles in the credential trust triangle

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 consented 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 dashed line in Figure 8 reflects the fact that a credential can only be used for transitive trust if the verifier has some degree of trust in the issuer and the policies, processes, and technology that the issuer uses to issue, deliver, and expire the credential.

Credentials have become such a routine part of the modern world that it is easy to overlook why they are so effective at conveying transitive trust. It is because they embody several principles:

1. **Trust is relational.** As Figure 8 shows, every credential is issued from a specific issuer to a specific holder in a specific relationship (citizen, employee, customer, graduate, etc.) That relationship is inherent to what the holder is able to prove to a verifier.
2. **Trust is directional.** A credential represents a one-way trust relationship from the issuer to the holder. Likewise, a credential presented by a holder to a verifier can be used to create a one-way trust relationship from the verifier to the holder. But if trust needs to run in the other directions—if the holder needs to trust either the issuer or the verifier—then the roles need to be reversed.
3. **Trust is contextual.** First, *every credential is issued to a holder in a specific context* (see #1 above). While there are many standardized credentials—passports, driving licenses, credit cards—there are also many specialized credentials that are only used in specific countries, industries, companies, or locales. Second, *every credential is presented by a holder to a verifier in a specific context*. This is where the flexibility and adaptability of credentials really shines. The variety of credentials that a verifier might choose to accept—and a holder might consent to share—together with the spectrum of different attributes or attestations those credentials may convey—give verifiers almost unlimited flexibility for how to arrive at a specific trust decision. Third, *every credential has an associated level of assurance.*[8] For example, when we rely on credentials to reflect the identity of the holder, the corresponding identity proofing process determines how much trust (or assurance) one can have in the credential. A library card is likely to have a lower level of assurance than a national ID card, for example.

The Governance Trust Diamond

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 to be widely adopted it must not only be easy to obtain, easy to use, and broadly interoperable—it must also be trusted by a large population of verifiers. This is precisely the kind of *transitive trust scaling problem* that we described at the start of this paper.

In the real world, this scaling problem has been solved very successfully by taking the credential trust triangle and adding a second trust triangle called the *governance trust triangle*. The combin-ation of these two trust triangles produces the *governance trust diamond* shown. in Figure 9.

[8] See ISO/IEC 29115 — https://www.iso.org/standard/45138.html

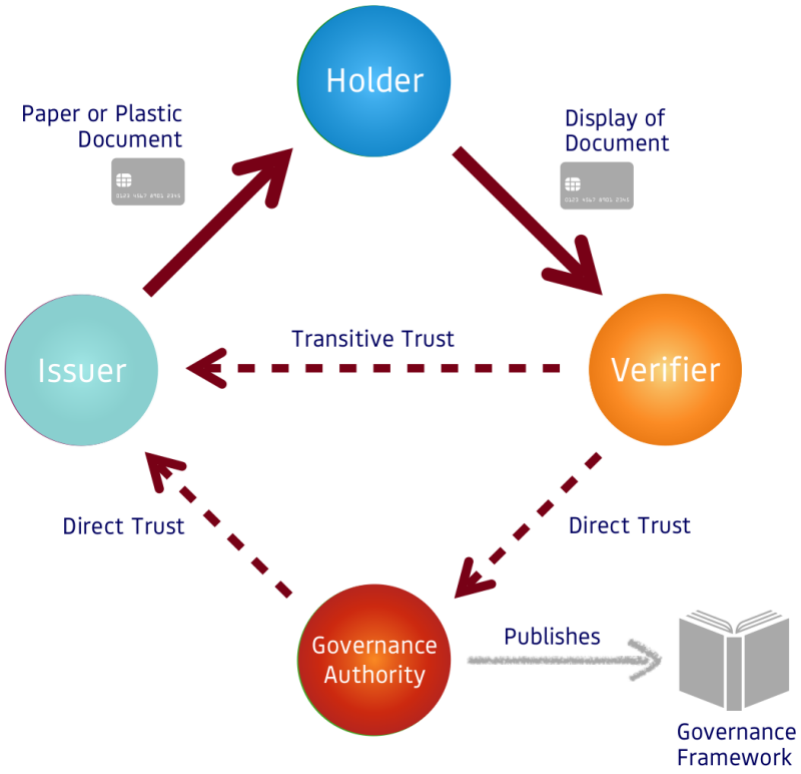

Figure 9. The governance trust diamond

The term "governing authority" (the green circle in Figure 9) does not mean this role is reserved for governments. To the contrary, any set of stakeholders who want to standardize the business, legal, and technical policies for issuing, holding, and verifying a set of credentials can serve as a governing authority. Furthermore, a governing authority can take any legal form—corporation, consortia, cooperative, informal community—but the purpose is always the same: develop and publish a *governance framework* that documents the policies and rules which the members of a trust community agree to follow in order to achieve their mutual trust objectives.

Perhaps the best-known examples of using the governance trust diamond to scale transitive trust are the credit card networks. Mastercard and Visa serve as the governing authorities for two of the world's largest payments networks, where the issuers are financial institutions all over the world, the holders are consumers everywhere, and the verifiers are tens of millions of merchants. The governance trust diamond for the Mastercard network is shown in Figure 10.[9]

[9] Mastercard is a founding Steering Member of the ToIP Foundation.

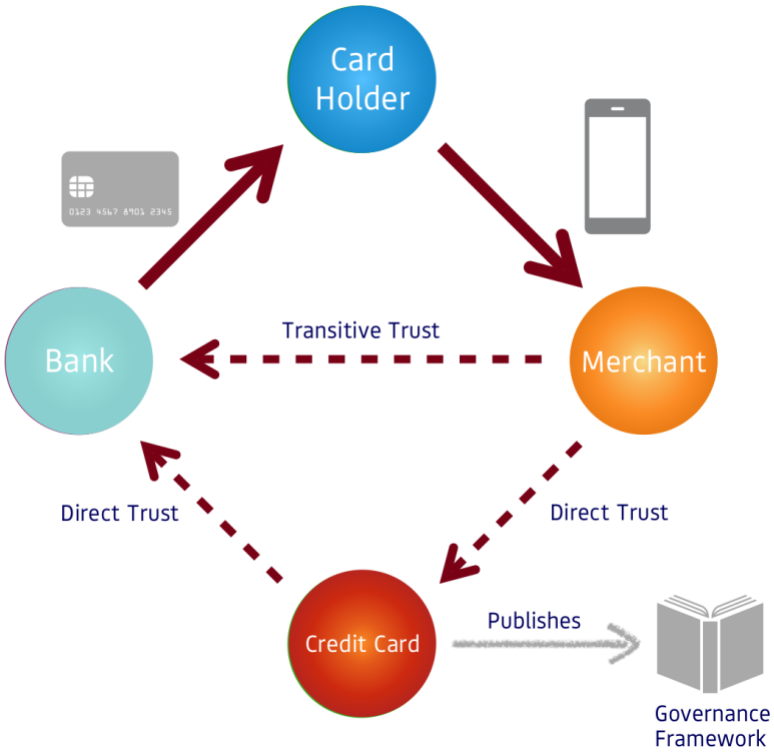

Figure 10. The governance trust diamond for the Mastercard payment network

Of course, the governance trust diamond applies to government-issued credentials of all kinds, from ID cards to health insurance cards to business licenses and countless other examples, as shown in Figure 11.

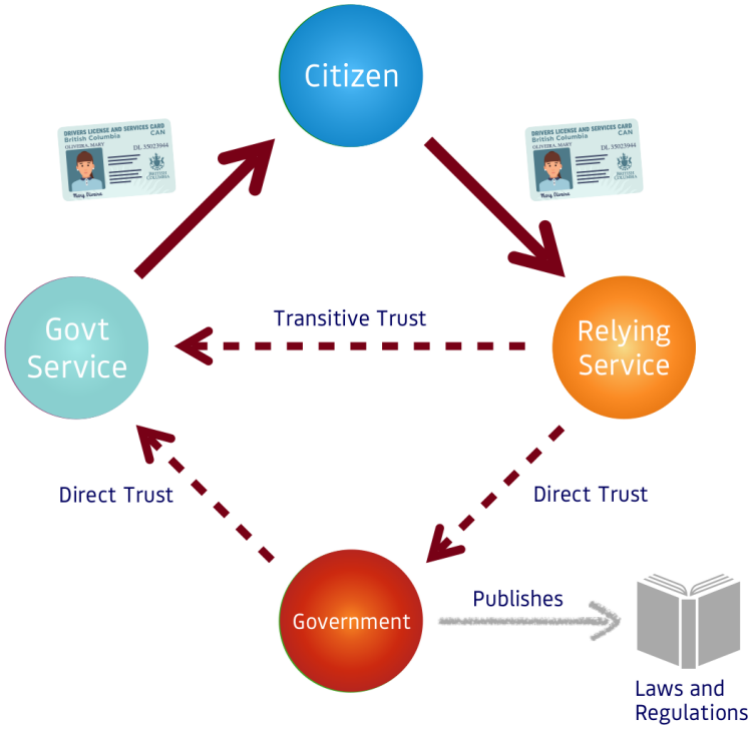

Figure 11. The governance trust diamond for government-issued credentials

The ultimate example of a governance framework for government-issued credentials is the standards for international passports governed by the International Civil Aviation Organization (ICAO), which includes documents such as Doc 9303 Machine Readable Travel Documents, the technical standard for machine-readable passports.

It should be clear that the governance trust diamond can apply to any type of credential issued within any trust community, large or small. One historical example is automobile association membership cards such as those issued by local chapters of the American Automobile Association or the Australian Automobile Association. The universal need for automobile roadside assistance means these cards are now accepted by tens of thousands of towing companies and other automobile service providers—as well as millions of hotels and other merchants offering discounts.

Applying This Model to the Digital World

Digital Credentials

Despite all the advantages, there are good reasons why the Internet didn't immediately adapt our real-world trust model of physical credentials to the digital world. Physical credentials are 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 because most digital files are trivial to copy—in milliseconds—unless they are protected by some form of cryptography, e.g., a digital signature, that proves their authenticity.

There are two basic approaches to providing such protection. One is to issue digital credentials to specialty devices designed to have the necessary security, privacy, and portability features. One example is smart cards—credit-card sized devices with low-power computer chips that can be read by specialized readers. Thousands of different smart card applications are in use today by many millions of holders for public transit, school ID, employee ID, etc.

The other approach is to issue digital credentials to standard computing devices (e.g., smartphones, tablets, laptops) and protect them using some form of public key infrastructure (PKI) to establish a chain of trust between issuers of digital certificates. These certificates are issued to holders who can present them to any verifier who trusts the PKI under which they were issued. This approach works entirely over any digital network—the Internet, local networks, mesh networks, or device-to-device connections such as Bluetooth, NFC, or QR codes.

Both approaches were simply too challenging to impose when the Internet was just getting on its feet. However, now that it is maturing, the benefits of introducing digital credentials and digital wallets are enormous. Imagine how much simpler the journey would be for a business owner like Sally, shown in Figure 12. 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.

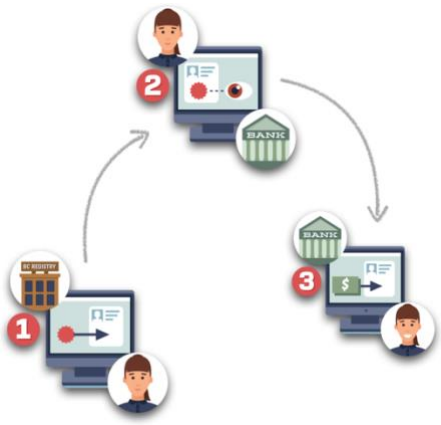

Figure 12. Digital credentials can simplify the journey of Sally, a business owner, in obtaining a bank account and license for her business

The Verifiable Credential Trust Triangle

Thankfully the promise of digital credentials was recognized several years ago by an active community within the World Wide Web Consortium (W3C).[10] They began the effort to standardize the file formats and digital signatures that would be needed for broad adoption. The result was the Verifiable Credentials Data Model 1.0 specification, which was approved as a full W3C standard in September 2019. Figure 13 illustrates the four basic steps in the sequence of issuing, holding, and presenting a verifiable credential.

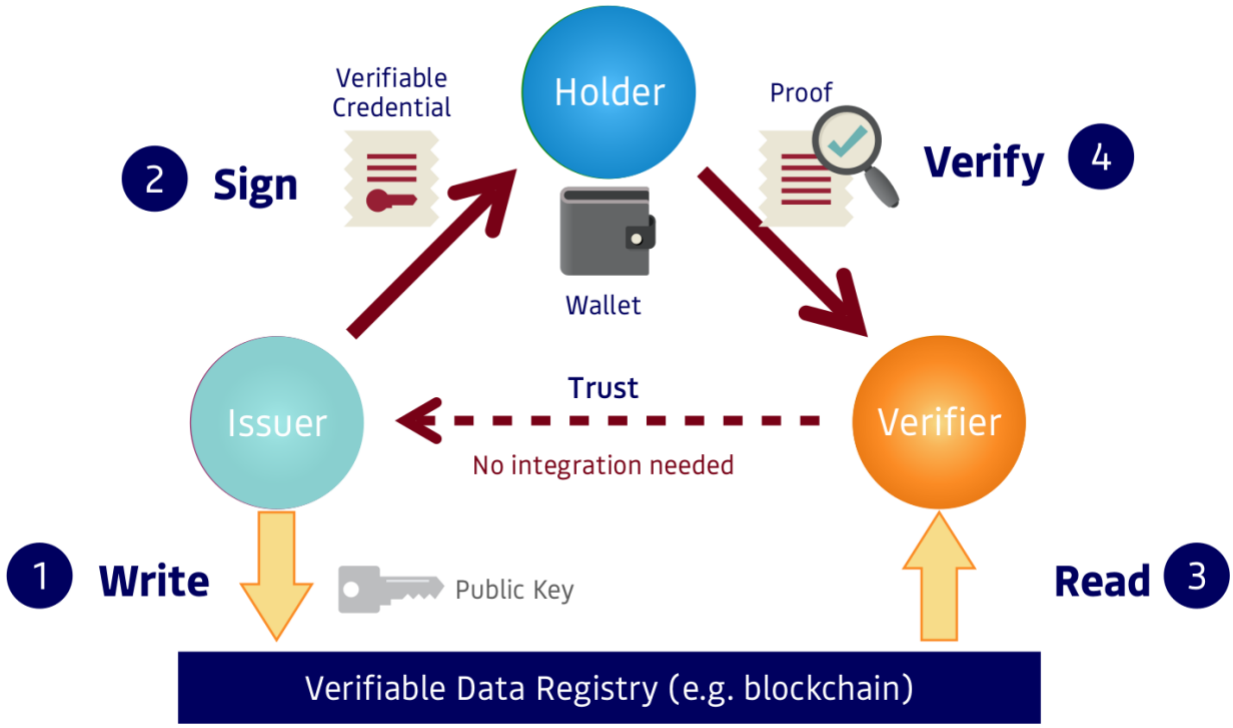

Figure 13. The four basic steps in the verifiable credential trust triangle

1. First the issuer writes a decentralized identifier (DID)[11] together with its public key (and any other cryptographic material needed to verify the issuer's credentials)[12] to some type of verifiable data registry (for example, a blockchain, a distributed database, or any other sufficiently trusted public utility accessible to verifiers).[13]

[10] This community is now the W3C Credentials Community Group, which continues to be a hotbed of innovation for the field of digital credentials.

[11] Final standardization of the W3C Decentralized Identifiers (DIDs) 1.0 specification is expected in Q1 2022.

[12] For an overview of the basics of public/private key cryptography, see https://en.wikipedia.org/wiki/Public-key_cryptography.

[13] Strictly speaking a DID is not required to issue a W3C-compliant verifiable credential; the issuer can also use conventional PKI-based digital certificates.

2. Second, the issuer uses its private key to digitally sign a verifiable credential and issues it to a qualified holder to store in his/her/its digital wallet. Note that to preserve privacy, this issuance process *does not need to involve any interaction with a verifiable data registry*—in other words, no personal data needs to be written to a blockchain or third-party data repository. The process can be fully confidential between the issuer and holder.
3. Third, a verifier requests a digital proof[14] 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 can use it to read the issuer's public key and other cryptographic data from the verifiable data registry.
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.

There are four important advantages to this process:

1. **No integration is needed between issuers and verifiers.** This is perhaps the single most important difference between the decentralized model and the federated model. The latter requires verifiers to communicate directly with issuers—an imposition that adds cost, complexity, availability and scaling issues, and—above all—privacy concerns about "phoning home" every time a credential is verified.
2. **It works for any type of digital credential.** Whereas federated identity models are typically limited to sharing of a narrow set of identity attributes, digital credentials enable issuers and holders to share any type of digitally signed data that can help a verifier make a trust decision.
3. **The process is actually claims-based, not credential-based.** Every digital credential is a collection of *claims*—digital assertions of any kind about the subject of the credential. For a digital driving license, for example, these claims might include name, address, birthdate, eye color, and so on. By standardizing how claims are packaged into digital credentials, and how verifiers can query for the claims they need, this claims-based architecture enables much larger, richer, and more fluid sets of contextually-relevant information to be shared between holders and verifiers.
4. **The process can be much more privacy-preserving than physical credentials.** When a holder needs to share a physical credential, such as a passport, driving license, or credit card, typically the verifier sees the entire credential—all of its claims—whether needed or not. This is both a security and a privacy issue. By contrast, digital credentials can use special types of cryptography—such as zero-knowledge proofs—that enable digital wallets to support selective disclosure of claims across any set of digital credentials. This enables verifiers to request and receive only the proof(s) they need to make a specific trust decision.

[14] For more about cryptographic proofs, see https://en.wikipedia.org/wiki/Provable_security.

The Digital Governance Trust Diamond

With verifiable credentials and digital wallets, we can use the same trust model—and conceptual 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 digital trust ecosystem*. This is the **digital governance trust diamond** in Figure 14.

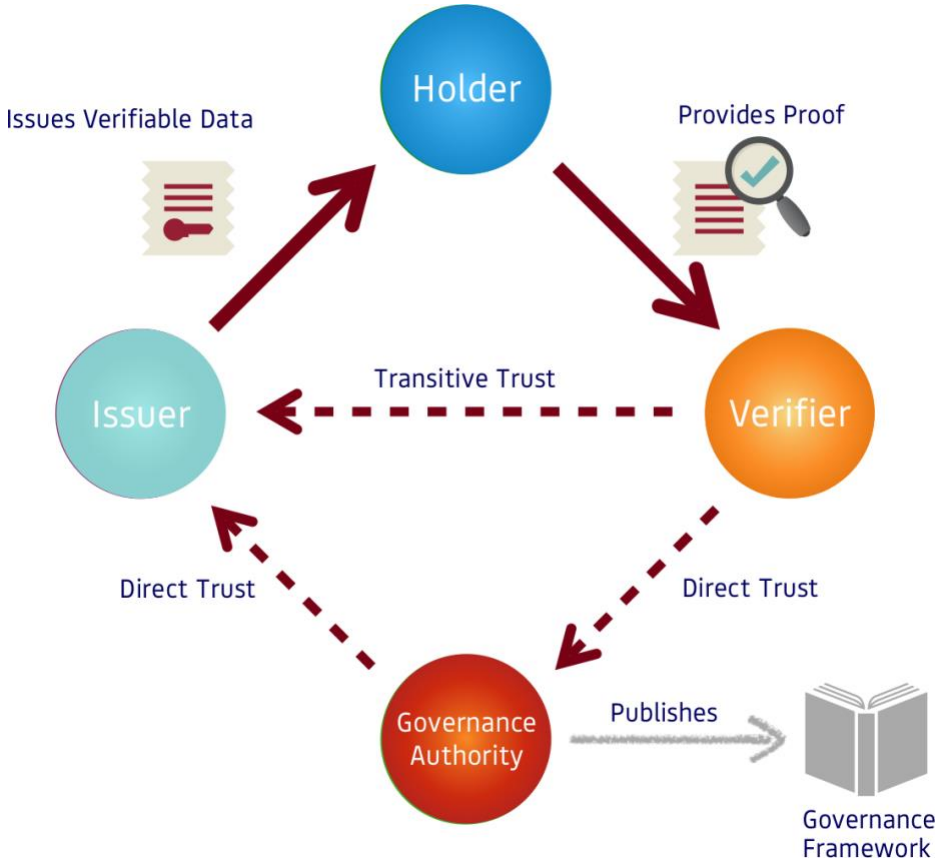

Figure 14. The digital governance trust diamond

As Figure 14 suggests, digital governance frameworks are the non-technical backbone of this new era of digital trust. Every digital credential in your wallet should be backed by a governance framework—large or small—that spells out the business, legal, and technical policies and rules under which that credential operates. By combining the technical trust that can be achieved with DIDs and verifiable credentials with the human trust codified in these governance frameworks, we can finally usher in a new era of Internet-scale decentralized digital trust infrastructure.

The ToIP Stack

The term “Trust Over IP” was coined by John Jordan, Executive Director of the Province of British Columbia Digital Trust Service, to describe a vision that he and other architects working on DIDs and digital credentials shared: that a “trust layer” for the Internet could be achieved by following the same architecture as the Internet: each peer running an instance of a standard “stack” of protocols just as each device on the Internet runs an instance of the TCP/IP stack.

This combination of technology and governance resulted in the two-sided, four-layer stack shown in Figure 15. Note that the lower two levels focus on meeting the technical requirements of digital trust, while the top two layers focus on meeting the human requirements.

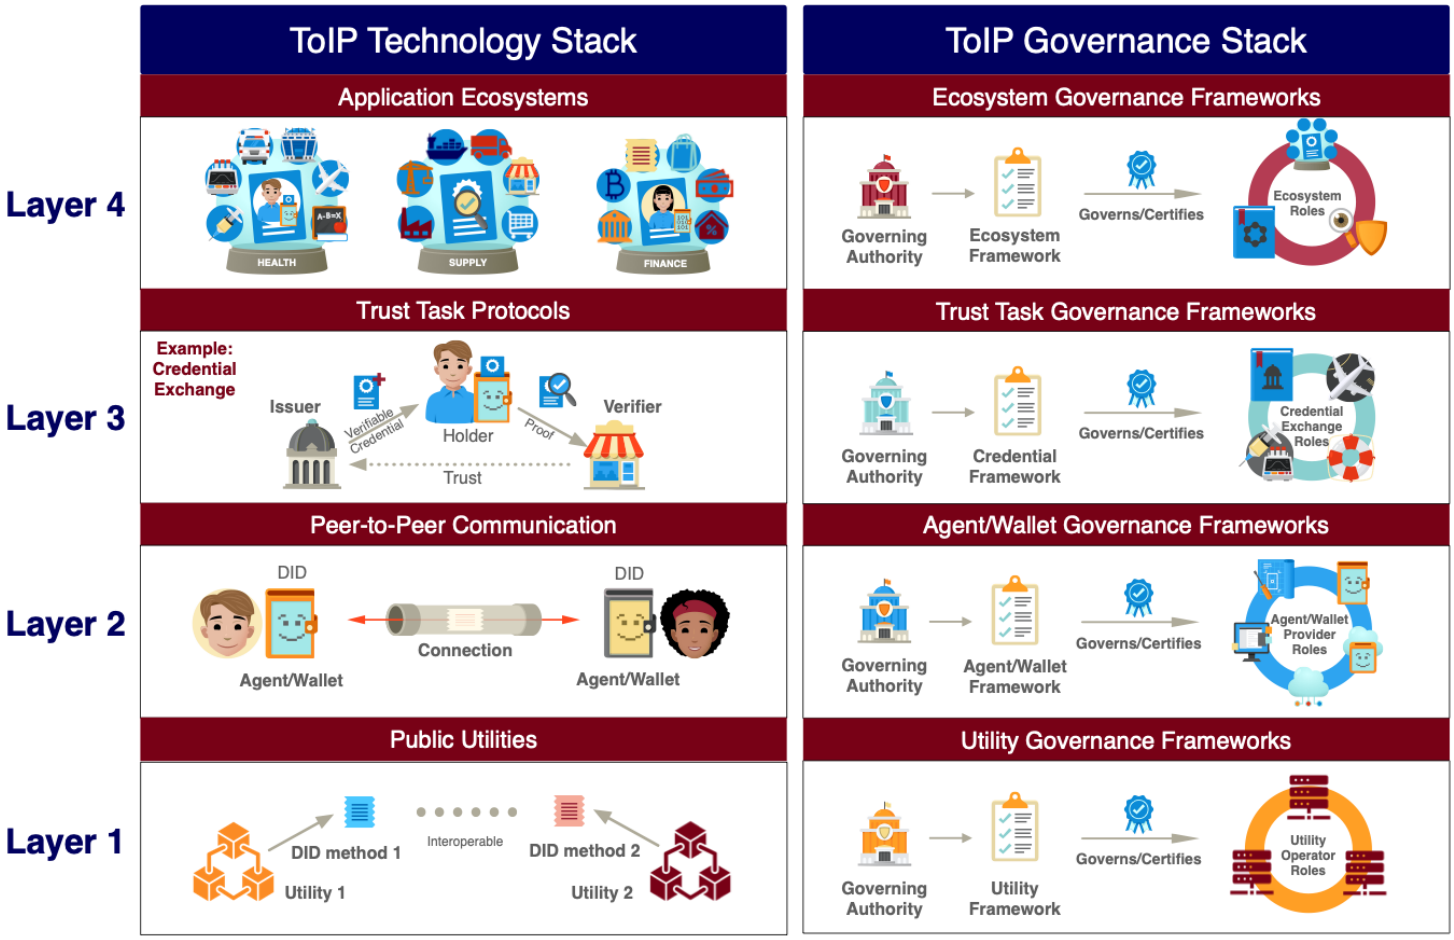

Figure 15. An overview of the four layers and two halves of the ToIP stack

What are the Design Principles Underlying this New Model?

The original designers of the ToIP stack realized this new model for decentralized digital trust infrastructure should follow two sets of *design principles*—principles that inform, guide, and constrain the design of a product, service, or system.

1. **Principles of computer network architecture**—above all these are the design principles underlying the Internet itself, the most important of which are articulated in IETF RFC 1958, *Architectural Principles of the Internet*.
2. **Principles of human network architecture**—above all the fundamental nature of how direct trust and transitive trust relationships are formed, strengthened, and (when necessary) broken.

When the ToIP Foundation was officially launched in May 2020, it was agreed that one of the first deliverables should be an articulation of these design principles. However, the process of exploring and refining these principles required an ongoing dialog between all four original (now eight) ToIP Working Groups. So the resulting document, Design Principles for the ToIP Stack, was not completed until December 2021. This document is highly recommended for anyone who wants to understand the fundamental principles guiding the design.

Why Four Layers?

All large-scale network infrastructure of any kind—physical or digital—follows the same basic four-layer model. Figure 16 illustrates how this model applies to our land transportation network.

Layer	Purpose	Tooling Required at each Layer
4	Transportation Industry (Market Applications)	
3	Signage & Traffic Controls (Rules of the Road)	
2	Cars & Trucks (Private Equipment)	
1	Roads & Highways (Public Utilities)	

Figure 16. The four basic layers of infrastructure for our land transportation network

Each layer solves a specific set of problems required to support the higher layers. For example, with our land and air transportation networks:

1. **Layer 1 is for the public utilities** that must be developed at great public investment to enable all the higher layers. With our land transportation network, these are the roads, highways, bridges, and tunnels that cost billions of dollars. With our air transportation network, these are the airports that must be constructed.
2. **Layer 2 is for the private equipment** that must be built to use the public utilities. For land transportation, these are the cars, trucks, motorcycles, and other vehicles we drive. For air transportation, these are the planes we fly.
3. **Layer 3 is for the "rules of the road"** needed to safely coordinate shared use of the infrastructure by an unlimited number of independent actors. For land transportation, these are our regulations, signs, and traffic signals. For air transportation, these are our regulations, air traffic control systems, GPS systems, and radio communication systems.
4. **Layer 4 is for the market applications** that need to be built on top of the first three layers to deliver value to end-users. For land transportation, these are taxis, buses, freight companies, car-rental services, ride-sharing services, and so on. For air transportation, these are airlines, charter services, travel services, etc.

With the ToIP stack, this is how the four layers apply:

1. **Layer 1 is for the public DID utilities** that are needed to look up and verify the current public keys of issuers of digital credentials. In public/private key infrastructure, these cryptographic "starting points" are called *roots of trust* or trust anchors.
2. **Layer 2 is for the private digital wallets and agents** needed by individuals, organizations, and digital "things" (or the digital twins of non-digital things) in order to accept, store, and exchange digital credentials over a standard peer-to-peer protocol such as DIDComm.
3. **Layer 3 is for the verifiable credential trust triangle** that enables establishment of transitive trust relationships between any three parties anywhere in the world using the data exchange formats and protocols for verifiable credentials.
4. **Layer 4 is for the market applications** needed to build healthy, vibrant digital trust ecosystems on top of this new decentralized digital trust infrastructure.

Why Two Halves?

Technology by itself is never sufficient to produce trust—simply because trust is a psychological belief of humans (or groups of humans). Therefore, technology must be harnessed to human behaviour to produce end-to-end trust. This is the role of governance—and the reason for the governance trust diamond at Layer 3 of the ToIP stack.

The insight that ultimately resulted in the two halves of the ToIP stack was the realization that governance also applies *at every other layer*. Governance of some kind—formal or informal, computer code or legal code—is required to drive business, legal, and social acceptance. For this reason, governing authorities (of some kind) and governance frameworks (of some kind) are needed at all four layers.

1. **At Layer 1, utility governance frameworks** are needed to establish trust in the policies and procedures used to operate public DID utilities (e.g., blockchains and distributed ledgers). Note that these governance frameworks may vary significantly depending on the cryptographic architecture of the utility. For example, public permissionless blockchain networks like Bitcoin or Ethereum use algorithmically-driven governance compared to public permissioned blockchain networks like Sovrin or IDunion.
2. **At Layer 2, wallet/agent governance frameworks** are needed to establish security, privacy, data protection, and interoperability standards for digital wallets and the digital agents that communicate between them.
3. **At Layer 3**, **credential governance frameworks** are needed to implement the governance trust diamond so verifiers have all the information they need to make trust decisions based on the verifiable credential proofs they are presented.
4. **At Layer 4, ecosystem governance frameworks** are needed to establish the policies and rules that will enable operation of entire digital trust ecosystems across all three lower layers. In many cases this will include recognizing independent governing authorities and governance frameworks that are authorized and/or supported at the lower layers. Ecosystem governance frameworks may also specify trust marks, trust registries, usability requirements, certification programs, and other mechanisms necessary to ensure the integrity and health of the entire ecosystem.

The ToIP Foundation

The Trust Over IP (ToIP) Foundation was launched in May 2020 with 27 original founding member organizations. Its vision for decentralized digital trust infrastructure was articulated in a Linux Foundation paper called The ToIP Stack published in August 2019 and subsequently turned into a December 2019 article in a special edition of IEEE Communications Standards Magazine on decentralized digital identity.

In its first year, the ToIP Foundation grew rapidly to over 300 member organizations and individuals. It also doubled from four to eight Working Groups, with over a dozen deliverables slated for release by the end of 2021. This surge of interest in decentralized digital trust infrastructure shows no signs of abating, particularly after the June 2021 announcement of the EU Digital Identity Wallet initiative and Apple and Google's announcements that they would begin to accept digitally-signed credentials in their proprietary digital wallets.

Membership Structure and Governance

The ToIP Foundation is hosted by the Linux Foundation under its Joint Development Foundation legal structure. There are three basic membership classes—Contributor, General, and Steering. The work of the Foundation takes place in working groups, within which there are self-organized Task Forces focused on specific interests. All ToIP members regardless of membership class can participate in any ToIP Working Groups and Task Forces.

The Foundation is governed by the Steering Committee, which works by consensus. The Steering Committee has 25 available voting seats comprised of a maximum of 15 representatives from Steering Members with over 100 employees and a maximum of 10 from Steering Members with less than 100 employees. Steering Committee members are elected by the Steering Members for a three-year term, with one-third of the members cycling each year.

Meetings and Collaboration

The Steering Committee meets twice a month, alternating between plenary and special topic sessions. The monthly All Members Meeting rotates between two formats:

1. **Working Group Update Meetings** encourage cross-collaboration and assist new members in deciding where they would like to contribute and learn.
2. **Special Topic Meetings** are a "deep dive" into topics of interest to the full organization.

All Working Group and Task Force meetings are open to every member. Meetings are held across different time zones to support global participation. See the ToIP Calendar page for a complete schedule. All meetings are held in Zoom rooms and recorded for asynchronous viewing. We also collaborate in the ToIP Slack workspace (with channels for each Working Group and Task Force), on our Confluence wiki, and in collaborative spaces such as GitHub.

ToIP Working Groups

This is a summary of the ToIP Working Groups as of December 2021. Check the Working Group section of the ToIP website for announcements of new Working Groups.

Name	Link	Description
Stack Working Groups		
Technology Stack Working Group	TSWG	Define (directly or by reference) the technical standards, test suites, and interoperability certification standards for the ToIP Technology Stack.
Governance Stack Working Group	GSWG	Specify requirements, templates, guides, roles, and processes for the ToIP Governance Stack.
Foundry Working Groups		
Ecosystem Foundry Working Group	EFWG	Facilitate a community of practice for establishing new ToIP digital trust ecosystems, whether hosted at the Linux Foundation or external to it.
Utility Foundry Working Group	UFWG	Facilitate a community of practice among governing authorities, implementers, operators, and service providers for ToIP Layer 1 public utilities.
Layer-Independent Working Groups		
Concepts & Terminology Working Group	CTWG	Facilitate ToIP Working Groups to develop the concepts, terms, and glossaries needed to establish shared understanding of their projects and deliverables.
Inputs and Semantics Working Group	ISWG	Define specifications and best practices for cohesion and interoperability for authentic data entry (“inputs”) and deterministic data capture (“semantics”).
Human Experience Working Group	HXWG	Develop insights & practical resources to enable ToIP stakeholders to improve outcomes for end-users.
Ecosystem-Specific Working Groups		
Good Health Pass Working Group	GHPWG	Facilitate a community of practice among implementers, issuers, holders, verifiers, and governing authorities of digital health passes for COVID-19 health status.

How to Engage with the ToIP Foundation

As with all Linux Foundation projects, the ToIP Foundation is an open collaborative organization that welcomes new members and provides multiple ways to participate, learn, contribute, and evangelize.

Membership

Membership enables your organization to tackle issues and solve problems in digital trust infrastructure that are beyond the ability of any one organization, governmental jurisdiction, or project ecosystem to solve on their own. The benefits are lower costs, reduced fraud, improved customer experience, faster time to market, and greater interoperability.

The ToIP Foundation offers three levels of membership:

1. **Contributor Members** may join at no cost—either as organizations or individuals—and can engage in all ToIP Working Groups and Task Forces. Contributor Member company logos do not appear on the ToIP website and are not featured in our press releases or other marketing materials.
2. **General Members** pay a nominal fee to support ToIP Foundation activities. General Member logos are included on the ToIP Foundation website and have the opportunity to be featured in our press releases, events, and other marketing materials.
3. **Steering Members** provide increased financial support for the Foundation and participate in governance via the ToIP Steering Committee. Steering Members also serve on our Communications Committee and are featured most prominently on our website and outgoing marketing communications.

For more information about joining the ToIP Foundation, see our Membership page.

Evangelism, Education, and Implementation

Whether you are a member or not, we encourage you to spread the message of the importance of interoperable decentralized digital trust infrastructure. If you find a ToIP white paper or other output that you find interesting and useful, please share it on social media, tweet it out using **#TrustOverIP**, or write a blog post on the topic and reference https://trustoverip.org/.

Most importantly we wish to see the digital layer of trust implemented and used by individuals and organizations around the world. So please take our specifications, recommendations, guides, glossaries, and other tools and build your piece of the puzzle or your digital trust ecosystem. Then tell us about it—we would be happy to share it as an example to the world.

The Road Ahead

The Gartner Hype Cycle for Emerging Technologies in 2021 named "decentralized identity" as the technology at the peak of "inflated expectations". This level of interest was one reason Gartner's #1 theme for innovation in 2021 was "*Engineering trust*". To quote Gartner's summary of this trend:

> *For IT teams to effectively lead technology-enabled business transformation, they must engineer a trusted business core. Trust requires security and reliability but must also be built on working practices that are repeatable, proven, scalable and innovative. These practices establish a resilient core and foundation for IT to deliver business value.*

What follows the peak of Gartner's hype curve is the "trough of disillusionment". This is the period where the hard work of fully developing, standardizing, hardening, and scaling a new technology must be done.

This is precisely the work of the ToIP Foundation. We recognize that decentralized digital trust infrastructure will not be an overnight success. It will take time to reach its full potential. We expect adoption to grow steadily as standards mature, open source code proliferates, and public and private sectors demonstrate the value of enabling strong, durable, privacy-respecting digital trust relationships between people, businesses, governments, and digital things everywhere.

Thus, we wanted to conclude this paper by articulating our realistic aspirations looking forward in one-, three-, and five-year time frames from the publication of this paper in December 2021.

One Year Horizon

By the end of 2022, the ToIP Foundation aims to accomplish the following goals:

- **Technology:** Deliver first-generation technical specifications for all four layers of the ToIP stack.
- **Governance:** Deliver first-generation governance framework recommendations, templates, worksheets, and guides for all four layers of the ToIP stack.
- **Adoption:**
    - At least one million digital credentials issued within ToIP digital trust ecosystems.
    - At least 10 digital trust ecosystems operating under ToIP-compliant governance frameworks.
    - At least three governments operating projects based on implementing ToIP-based infrastructure.
    - All major IT analyst firms providing coverage of ToIP.

Three Year Horizon

By the end of 2025, our aim is to achieve the following:

- **Technology:**
  - Deliver second-generation technical specifications for all four layers of the ToIP stack.
  - Contribute one or more ToIP technical specifications to Standards Development Organizations (e.g., ISO, IETF, W3C) to become international standards.
- **Governance:**
  - Deliver second-generation governance framework recommendations, templates, worksheets, and guides for all four layers of the ToIP stack.
  - International standards for security, privacy, and data protection begin to reference ToIP-compliant governance frameworks.
- **Adoption:**
  - At least 50 million digital credentials issued within ToIP digital trust ecosystems.
  - At least 100 digital trust ecosystems operating under ToIP-compliant governance frameworks.
  - At least 10 governments operating projects based on implementing ToIP-based infrastructure.
  - One or more industry events, conferences, or magazines devoted to ToIP.

Five Year Horizon

By the end of 2027, our goal is to have catalysed the following changes in the market:

1. The ToIP stack is beginning to be shipped with OEM devices.
2. Support for the ToIP stack is built into at least one major browser.
3. App stores have added labels for vendors to indicate whether a product is ToIP-compatible.
4. Governments have begun passing legislation or regulations requiring the use of ToIP-based digital trust infrastructure.

We invite you to join us on this journey to a safer and more trusted digital world.

TRUST Over IP FOUNDATION

The Trust Over IP Foundation (ToIP) is hosted by the Linux Foundation under its Joint Development Foundation legal structure. We produce a wide range of tools and deliverables organized into five categories:

- Specifications to be implemented in code
- Recommendations to be followed in practice
- Guides to be executed in operation
- White Papers to assist in decision making
- Glossaries to be incorporated in other documents

ToIP is a membership organization with three classes—Contributor, General, and Steering.

The work of the Foundation all takes place in Working Groups, within which there are Task Forces self-organized around specific interests. All ToIP members regardless of membership class may participate in all ToIP Working Groups and Task Forces.

When you join ToIP, you are joining a community of individuals and organizations committed to solving the toughest technical and human centric problems of digital trust. Your involvement will shape the future of how trust is managed across the Internet, in commerce, and throughout our digital lives. The benefits of joining our collaborative community are that together we can tackle issues that no single organization, governmental jurisdiction, or project ecosystem can solve by themselves. The results are lower costs for security, privacy, and compliance; dramatically improved customer experience, accelerated digital transformation, and simplified cross-system integration.

To learn more about the Trust Over IP Foundation please visit our website, https://trustoverip.org.

Licensing Information:
All Trust Over IP Foundation deliverables are published under the following licenses:

Copyright mode: Creative Commons Attribution 4.0 International licenses
http://creativecommons.org/licenses/by/4.0/legalcode

Patent mode: W3C Mode (based on the W3C Patent Policy)
http://www.w3.org/Consortium/Patent-Policy-20040205

Source code: Apache 2.0.
http://www.apache.org/licenses/LICENSE-2.0.htm