TRUST Over IP FOUNDATION

Evolution of the ToIP Stack

A Reference Document from the ToIP Foundation

https://trustoverip.org/

Version 1.0 — 14 November 2022

***This is a living document** — the latest version can be found on this web page:*

https://trustoverip.org/our-work/evolution-of-the-toip-stack/

Table of Contents

Status of This Document

This is a living document. The first version was published by the ToIP Foundation on 14 November 2022 in conjunction with the announcement of the first public review draft of the ToIP Technology Architecture Specification V1. The Foundation plans to publish an updated version at each major milestone in the evolution of the ToIP stack as described in this document. Each new version will cover the progress along the four evolutionary stages and include an updated "map" in the Mapping Existing Technologies to the ToIP Stack section.

Introduction

The ToIP stack was first envisioned in 2019 as a new approach to Internet-scale decentralized digital trust infrastructure. It was inspired by a combination of industry developments:

1. The approval of W3C Verifiable Credentials 1.0 as the first global standard for digitally-signed credentials of all kinds (driver's licenses, passports, health cards, employment cards, credit cards, diplomas, tickets, and so on).
2. The emergence of distributed ledger technology as a source of strong, immutable, decentralized roots of trust for decentralized identifiers (DIDs) and cryptographic keys.
3. The growing ubiquity of mobile phones and cloud computing, making digital communications of all types more affordable and accessible.

As these technologies were applied to longstanding challenges in digital trust, a new Internet-scale digital identity model began to emerge that has gone by a succession of names:

- Blockchain identity
- Self-sovereign identity (SSI)
- Decentralized identity

As this new model matured, a group of architects working in the space recognized the model was developing into a four-layer “stack” for Internet-scale digital trust that had strong parallels with the four-layer TCP/IP stack whose adoption fifty years earlier gave us a pivotal technology that empowers the Internet as we know it today. The architects felt that an industry-wide effort to fully define and standardize this stack could do for globally-interoperable trust networking what the TCP/IP stack had done for data networking. This culminated in the launch of the Trust Over IP (ToIP) Foundation in May 2020.

The ToIP Stack

Figure 1 is the conceptual diagram of the ToIP stack used by the ToIP Foundation.

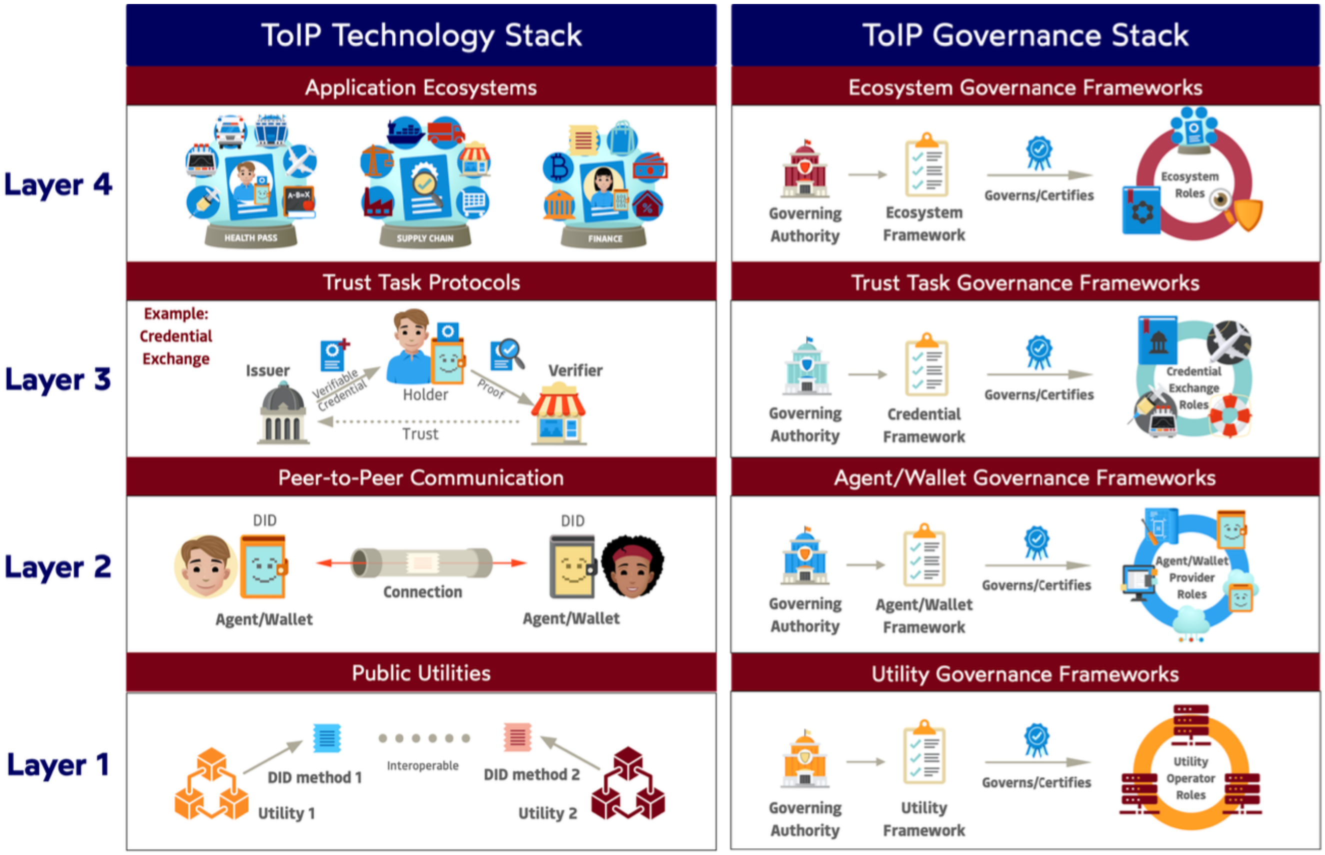

Figure 1: Conceptual diagram of the ToIP stack

Perhaps the most striking aspect of this diagram is that it is not a single “stack” but consists of two parallel, interlocking stacks, one for technology and one for governance. While the TCP/IP stack had certain types of governance, such as how it handled network segmentation, inter/intra domain routing, domain resolution and other functions, those types of governance has not proved to be sufficient to deal with the real-world trust issues that appeared after its exponential growth. The need for decentralized governance of decentralized digital trust infrastructure was one of the core insights that led to the founding of the ToIP Foundation to fully define both sides of the stack.

The ToIP Technology Stack

Although the technical requirements for each of the four layers of the ToIP Technology Stack are now available in the public review draft of the ToIP Technology Architecture Specification V1.0 (GitHub version, PDF version), at a very high level they can be summarized as:

1. **Layer 1: Trust Support**. This first layer contains the foundational components necessary to support the higher layers, e.g., trusted computing modules, secure storage, network transport protocols, and external verifiable data registries (to provide decentralized roots of trust that can span within and across different digital trust ecosystems).
2. **Layer 2: Trust Spanning**. This is the layer that enables the establishment of a trusted connection (ephemeral or persistent) between any two peers using a single standard trust spanning protocol. *This layer is to the ToIP stack what the IP layer is to the TCP/IP stack.*
3. **Layer 3: Trust Tasks**. A trust task is an element, or component, of an application that consists of one or more well-defined atomic actions that support the overall trust objectives of the application. Verifiable credential issuance and presentation, payments and other forms of value exchange, reservations, auction bids—all are examples of trust tasks. Just as many higher level protocols are built on top of the IP spanning protocol in the TCP/IP stack, many higher level trust task protocols can be built on top of the ToIP trust spanning protocol.
4. **Layer 4: Trust Applications**. Just as all Internet-enabled applications ultimately ride on top of the TCP/IP stack, all ToIP-enable applications operate at Layer 4 of the ToIP stack. A ToIP-enabled application, executing on a ToIP endpoint, calls Layer 3 trust task protocols or the Layer 2 trust spanning protocol directly in order to form trusted connections with other ToIP endpoints. Once those connections have been established the applications may engage in trusted interactions and execute trust tasks within one or more digital trust ecosystems.

The ToIP Governance Stack

It is a platitude in cybersecurity that “technology alone is necessary but not sufficient to achieve trust”—in particular to achieve *legal, business, or social trust*. The latter requires binding technology to the policies and rules that people, organizations, and governments agree to follow using that technology.

Such policies and rules exist at all four layers of the stack, which is why the ToIP Governance Stack recommends **governance frameworks** (aka **trust frameworks**) at all four layers:

1. **Layer 1: Trust Support**. Trust at any higher layer is only as strong as the parties' confidence in any shared trust roots or trust anchors. In today’s X.509-based Web trust infrastructure, this means confidence in the roughly 200 certification authorities (CAs) around the world. In tomorrow’s decentralized digital trust infrastructure, it means confidence in the governance frameworks for both: a) the trusted computing and secure

storage elements of private digital wallets, and b) the utilities that serve as public or private *verifiable data registries* for storing decentralized identifiers (DIDs), DID documents, trust lists, and other shared sources of truth.

2. **Layer 2: Trust Spanning**. As the “active” component of a digital wallet, a digital agent that speaks the ToIP trust spanning protocol is second only to an operating system as the software component that most needs to be trusted by its controller. The EU, Canada, Bhutan, and other countries are already moving down the path of mandating governance requirements for digital agents and digital wallets. This is what enables trust to move out to the decentralized “edges” of the network.
3. **Layer 3: Trust Tasks**. Every real-world trust task — from issuing a verifiable credential to initiating a payment to participating in an auction — requires actors (individuals or organizations) to behave by a known set of policies and rules. So, by definition, the level of trust the parties can achieve depends on their level of confidence in a shared governance framework (formal or informal, manual or algorithmic) for that trust task.
4. **Layer 4: Trust Applications**. At the application layer, governance frameworks can establish the rules for how one or more ToIP-enabled applications can interoperate and collaborate within and across digital trust ecosystems. Such ecosystems can be very narrow—such as within a single company, town or university—or they can be very broad, such as digital health or vaccination credentials that are recognized and accepted worldwide. Such governance frameworks can also mandate the use of specific user interface affordances for ToIP-enable applications, such as QR codes, trust marks, or accessibility requirements.

The ToIP Governance Stack Working Group published its first-generation ToIP Governance Architecture Specifications in January 2022 — see this ToIP blog post for an overview and links.

The Evolutionary Process

From the outset, the organizers of the ToIP Foundation knew that they needed to play “the long game”. Designing, building, testing, and adopting a decentralized trust layer for the Internet would take at least several years, if not a decade or more. But the benefits would be worth it.

The Four Stages

Two-and-a-half years into the work, the ToIP Foundation has developed a clear picture of the four stages required. These are summarized in Figure 2:

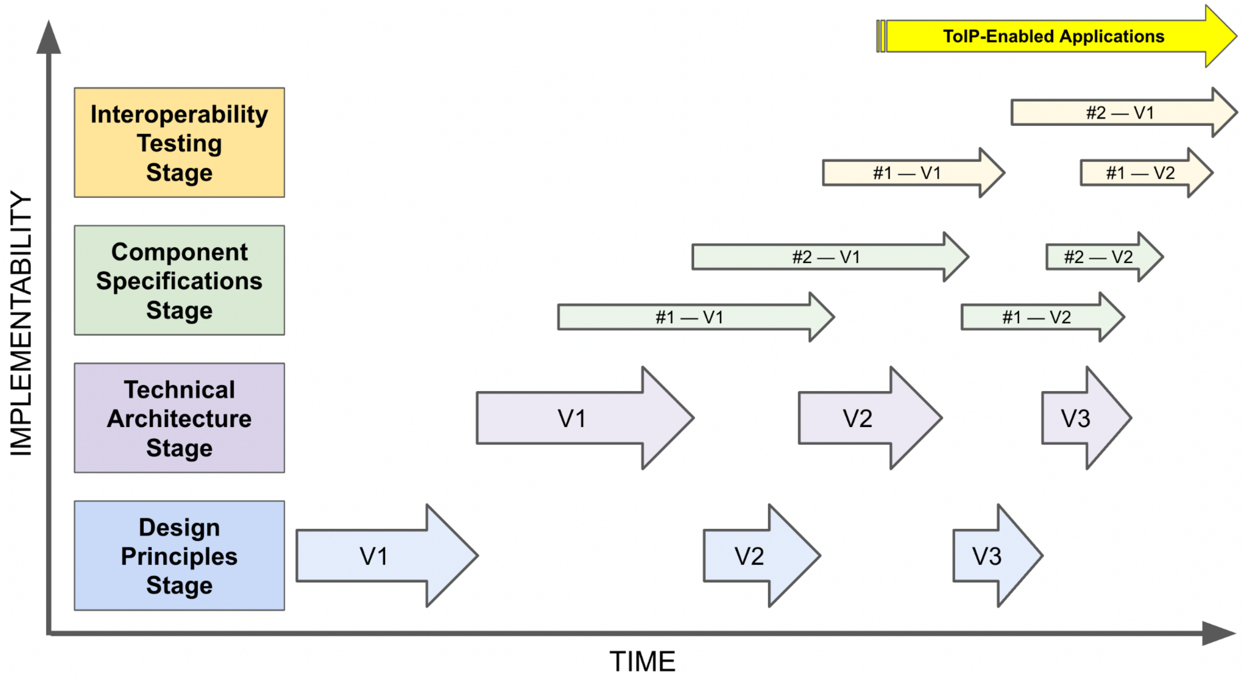
ToIP-Enabled Applications
Interoperability Testing Stage
#2 — V1
#1 — V1
#1 — V2
Component Specifications Stage
#2 — V1
#2 — V2
#1 — V1
#1 — V2
Technical Architecture Stage
V1
V2
V3
Design Principles Stage
V1
V2
V3
IMPLEMENTABILITY
TIME

Figure 2: The four evolutionary stages of development of the ToIP stack

Stage #1: Design Principles for the ToIP Stack

The need for the first stage was recognized when the Foundation’s membership swelled to over 200 organizations after the first year. Before we could dig deeply into the architecture of the ToIP stack, we needed to agree on the fundamental principles that would inform how the Foundation would achieve its stated objectives.

This work began in the summer of 2021. While we anticipated this would only take “a few sprints”, the work was so intensive and engaging that it continued for six months. The final document, Design Principles for the ToIP Stack V1.0, was approved by the ToIP Steering Committee on 17 November 2021.

We strongly recommend reading the 17 principles in this document for a complete perspective on the underlying design philosophy of ToIP architecture.

NOTE: As with the deliverables of each stage, we expect the Design Principles to further evolve as we learn more. For example, we are already planning a small revision based on insights gained in the development of the ToIP Technology Architecture Specification V1 (below). However we do not expect any major changes to the Design Principles, only small additions.

Stage #2: ToIP Technology Architecture Specification

Once the Design Principles were ready, the ToIP Technology Stack Working Group formed the Technology Architecture Task Force in November 2021 to begin work on the first version ToIP Technology Architecture Specification.

A little over a year later, the first public review draft of the ToIP Technology Architecture V1.0 Specification (GitHub version, PDF version) was published by the ToIP Foundation on 14 November 2022 for discussion at Internet Identity Workshop #35. Figure 3 graphically illustrates where this milestone lies in the evolutionary progression.

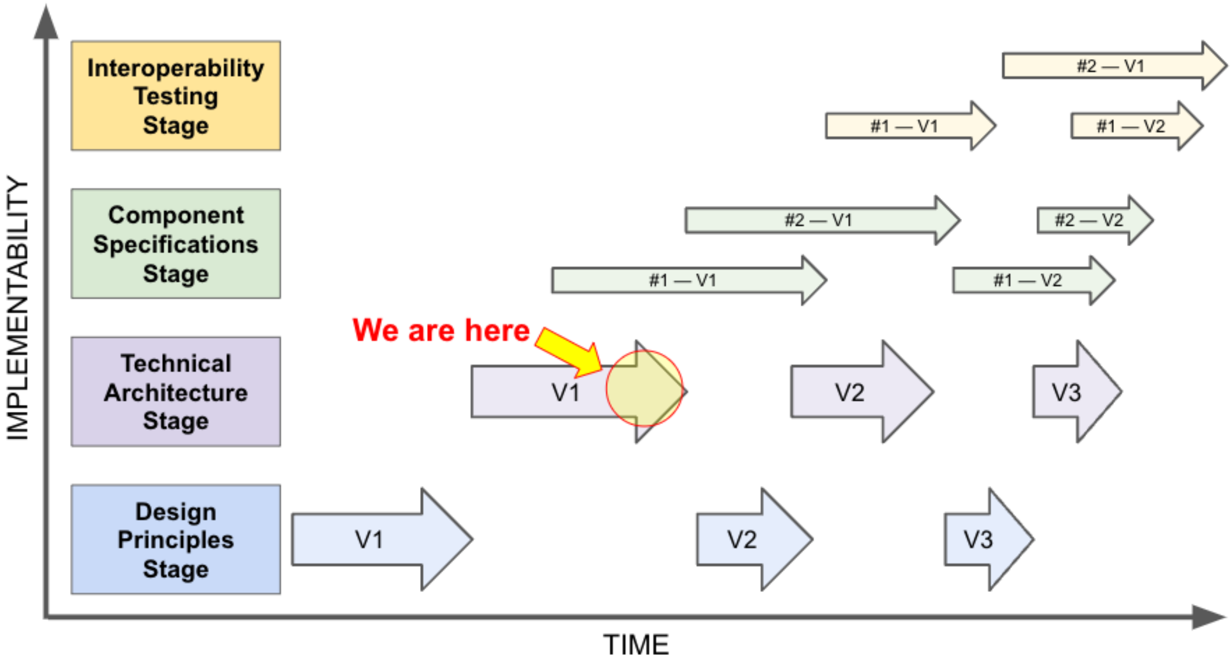

Figure 3: Where we currently are in the evolution of the ToIP stack

Stage #3: Component Specifications

The purpose of the ToIP Technology Architecture Specification is to enumerate the requirements that specific **component specifications** need to meet and inform the Interoperability work that must build upon these specifications. A component specification defines or profiles either a protocol, an interface, an affordance — or some combination of these — that accomplishes the design goals of one or more of the four layers of the ToIP Technology Stack.

A number of candidate component specifications already exist — see Mapping Existing Technologies to the ToIP Technology Stack. In this stage, a major thrust of ToIP Foundation members will be to help advance — and encourage others to help advance — these existing and other new component specifications to the point where a set can complete the ToIP stack "top-to-bottom". This milestone will trigger the final evolutionary stage necessary to begin producing ToIP-enabled applications.

Stage #4: Interoperability Profiles and Test Suites

To be used in production, multiple implementations of the ToIP stack must be proven to be interoperable to the satisfaction of all stakeholders within one or more digital trust ecosystems. Thus this stage will require:

1. Defining specific **interoperability targets** (roles or personas) for the various technology components of the ToIP stack.
2. Define how **ToIP interoperability profiles** (TIPs) can be written against these interoperability targets to specify the requirements of a particular digital trust ecosystem.
3. Specify how third parties, such as certification testing labs, can perform certification testing against implementations of each of these targets.
4. Specify how trust marks can be awarded upon certification.

When robust interoperability certification is available in the market, the first interoperability profiles of the ToIP stack will be ready for marketplace adoption.

Mapping Existing Technologies into the ToIP Technology Stack

With the first version of the ToIP Technology Architecture Specification (TAS) now entering public review, we can start the process of mapping where potential component specifications fit within the ToIP Technology Stack (and cross-tie into the ToIP Governance Stack). Each of these can be compared with the requirements specified in the TAS (see Appendix A of the TAS for a consolidated layer-by-layer list).

IMPORTANT NOTES:

1. *The following mappings are a work-in-progress that will initially be discussed in sessions of the Internet Identity Workshop #35, 15-17 November 2022. We expect the next version of this living document to contain updates and improvements based on those conversations as well as any other feedback received during the TAS public review period.*
2. *This mapping does not currently include an assessment of any potential component specification against the requirements of the TAS. That assessment will either be added later or published separately.*
3. ***These mappings are not exhaustive, nor are they intended to be an endorsement of any particular technology or specification.** Any normative adoption of component specifications into the ToIP stack as well as any interoperability profiles and test suites is future work that has not yet been defined.*
4. *Some technologies are listed at more than one level as their specifications currently define functionality that may belong at different layers of the ToIP stack.*

Layer 1

Specification	Home	Description
Decentralized Identifiers (DIDs) 1.0	W3C	DIDs are a new type of Uniform Resource Identifier (URI) that is cryptographically verifiable and discoverable without requiring a centralized registry.
Key Event Receipt Infrastructure (KERI) Witnesses	IETF (planned)	KERI defines a decentralized identification and key management infrastructure based on self-certifying Autonomic Identifiers (AIDs) and Autonomic Namespaces (ANs) as a primary root-of-trust.
X.509 Digital Certificates	ITU ISO/IEC	X.509 is a worldwide standard for PKI digital certificates from ISO (the International Organization for Standardization) and IEC (the International Electrotechnical Commission).

Layer 2

Specification	Home	Description
Decentralized Identifiers (DIDs) 1.0	W3C	DIDs are a type of Uniform Resource Identifier (URI) that is cryptographically verifiable and discoverable without requiring a centralized registry.
DIDComm 2.0	DIF (IETF planned)	DIDComm Messaging provides a secure, private communication infrastructure built atop the decentralized design of DIDs.
KERI with Composable Event Streaming Representation (CESR)	IETF (planned)	The Composable Event Streaming Representation (CESR) is dual text-binary encoding format that enables text-binary concatenation composability.
Decentralized Web Node (DWN)	DIF	A Decentralized Web Node (DWN) is a data storage and message relay mechanism that entities can use to locate public or private permissioned data related to a given Decentralized Identifier (DID).

Layer 3

Credential Exchange Protocols

Specification	Home	Description
W3C VC API	W3C	Verifiable credentials provide a mechanism to express credentials on the Web in a way that is cryptographically secure, privacy respecting, and machine-verifiable. This specification provides data models and HTTP protocols to issue, verify, present, and manage data used in such an ecosystem.
ISO mDL/mDOC	ISO	The ISO 18013-5 standard defines a mobile driver's license (mDL) that can replace a physical driver's license. A more general and extensible form of mDL is called mDOC.
OIDC4VC Issuance	OpenID Foundation	Defines an API and corresponding OAuth 2.0-based authorization mechanisms for the issuance of verifiable credentials.
OIDC4VP	OpenID Foundation	Defines an extension of OpenID Connect to allow presentation of claims in the form of W3C Verifiable Credentials as part of the OIDC protocol flow.
Hyperledger Aries Interop Profile	Hyper-ledger	Each Aries Interop Profile (AIP) version provides a set of Aries RFCs to enable interoperable exchange of verifiable credentials between Aries digital agents and wallets.

Credential Formats/Signature Suites

This area of potential component specifications is very complex. The mapping for this section will be based on (or link out to) the Credential Profile Comparison Matrix described in this paper.

Other Trust Tasks

Specification	Home	Description
Decentralized Web Node (DWN)	DIF	A Decentralized Web Node (DWN) is a data storage and message relay mechanism entities can use to locate public or private permissioned data related to a given Decentralized Identifier (DID).
Secure QR Codes	OASIS	Defines the use of QR Codes as a replacement for a username and password in user login authentication.

Layer 4

Specification	Home	Description
Secure QR Codes	OASIS	Defines the use of QR Codes as a replacement for a username and password in user login authentication.

How to Get Involved

Defining the ToIP stack is the mission of the ToIP Foundation, and in particular the focus of our Technology Stack Working Group and Governance Stack Working Group. The Foundation is also in the process of forming the **ToIP Interoperability Working Group** to develop interoperability profiles, test suites, certification frameworks, and other tools needed to achieve the full promise of the ToIP stack. We invite you to join the ToIP Foundation — either as an organization or as an individual — and participate in these Working Groups.

You can also contribute directly on the public review draft of the ToIP Technology Architecture Specification (Github version, PDF version) even if you are not a ToIP member — see the instructions on this page.