

# Trust over IP (ToIP) Technology Architecture Specification

---

## Revision History

- Public Review Draft 01 (PR1) — 14 November 2022

For the most recent version of this specification, please see [this ToIP web page](#).

## Editors

~ [Daniel Bachenheimer](#) (Accenture) ~ [Wenjing Chu](#) (Futurewei Technologies, Inc) ~ [Darrell O' Donnel](#) (Continuum Loop) ~ [Andor Kesselman](#) (Benri) ~ [Antti Kettunen](#) ~ [Christine Martin](#) (Continuum Loop) ~ [Drummond Reed](#) (Gen/Avast) ~ [Jo Spencer](#) (460degrees / Sezoo) ~ [Allan Thomson](#) (Gen/Avast)

## Contributors

~ [Jacques Bikoundou](#) ~ [Tim Bouma](#) (CIO Strategy Council) ~ [Kevin Dean](#) (GS1 Global Office) ~ [Judith Fleenor](#) (Trust Over IP Foundation) ~ [Sid Haniff](#) (Datasoc) ~ [Daniel Hardman](#) (Provenant) ~ [Isaac Henderson](#) ~ [John Jordan](#) (Province of British Columbia) ~ [Vikas Malhotra](#) (WOPLLI Technologies) ~ [Sankarshan Mukhopadhyay](#) (Dhiway Networks) ~ [Sumabala Nair](#) (IBM) ~ [Vinod Panicker](#) (Wipro) ~ [Scott Perry](#) (Schellman) ~ [Vladimir Simjanoski](#) (Blokverse) ~ [P. A. Subrahmanyam](#) (CyberKnowledge) ~ [Bart Suichies](#) ~ [Samuel Smith](#) (ProSapien LLC) ~ [Neil Thomson](#) (QueryVision) ~ [Alex Tweeddale](#) (cheqd) ~ [Mattia Zago](#) (Monokee) ~ [Vlad Zubenko](#) (ETS)

*Copyright: 2022 Trust Over IP Foundation*

## Table of Contents

- [1. Introduction](#)
- [2. Terminology](#)
- [3. Motivations](#)
- [4. Audience, Purpose and Scope](#)
- [5. Example Use Cases](#)

- 6. Reference Architecture Overview
  - 6.1 Design Goals
  - 6.2 The Four Layer Pattern
  - 6.3 High-Level System Architecture
  - 6.4 ToIP Identifiers
- 7. Endpoint Systems and the Layered Stack
  - 7.1 Endpoint Systems
  - 7.2 Layer 1: Trust Support
  - 7.3 Layer 2: Trust Spanning
  - 7.4 Layer 3: Trust Tasks
  - 7.5 Layer 4: Trust Applications
- 8. The ToIP Trust Spanning Protocol
  - 8.1 Overview
  - 8.2 Identifiers
  - 8.3 Messages
  - 8.4 Routing
  - 8.5 Interface to Layer 1
- 9. Intermediary Systems
- 10. Supporting Systems
  - 10.1 Overview
  - 10.2 Example 1 - A DID Method
  - 10.3 Example 2 - A KERI Witness
  - 10.4 Generalization
- 11. Endpoint System Interoperability
  - 11.1 Interoperability between Endpoint Systems Using Decentralized Identifiers
  - 11.2 Interoperability between Endpoints Using Other Verifiable Identifiers
- 12. Integration with the ToIP Governance Stack
- 13. References
- About the ToIP Foundation
- Appendix A: Consolidated Requirements
- Appendix B: Consolidated Views of the ToIP Technology Stack
  - Functional Hourglass View
  - Sphere-of-Influence View
  - Interaction Pattern View
- Appendix C: Mapping of Existing Technologies into the ToIP Technology Stack

# 1. Introduction

The mission of the [Trust over IP \(ToIP\) Foundation](#) is to define an overall architecture for Internet-scale digital trust that combines cryptographic assurance at the machine layers (technology) with human accountability at the business, legal, and social layers (governance). Together these two halves form a complete four-layer architecture for decentralized digital trust infrastructure known as the [ToIP stack](#). Figure 1 is a conceptual diagram of this stack:

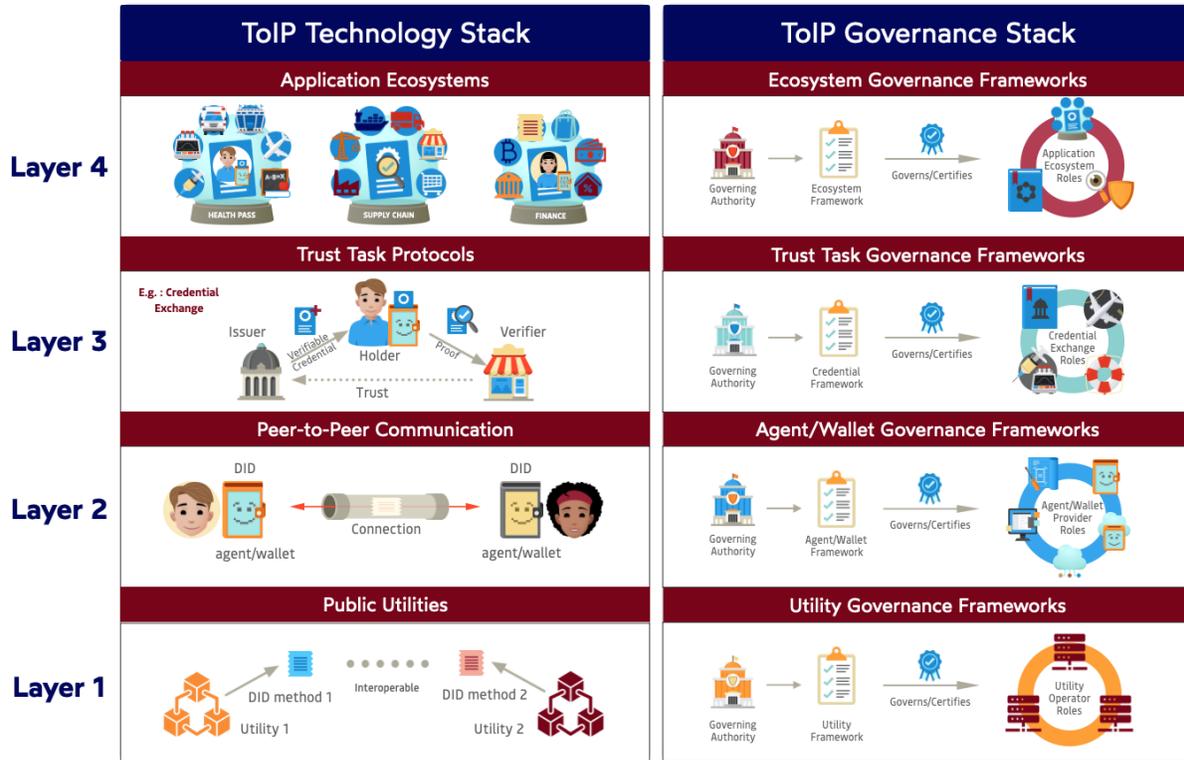


Figure 1: Conceptual diagram of the ToIP stack

This document is the normative specification for the high-level architecture of the ToIP technology stack (the left half of Figure 1). It is a deliverable of the [Technology Stack Working Group](#) at the ToIP Foundation. It is recommended to read this document in conjunction with these three other documents from the ToIP Foundation in the following order:

1. [Introduction to ToIP](#) is our white paper that provides an overall introduction to the emergence of decentralized digital trust infrastructure. It explains the origin and basic structure of the ToIP stack together with the mission and activities of the ToIP Foundation.

2. [Evolution of the ToIP Stack](#) is a companion document to this specification that explains the overall process the ToIP Foundation is following in the development of the ToIP stack. It is recommended for anyone seeking to understand how the work of the ToIP Foundation relates to that of adjacent non-profit organizations such as the Decentralized Identity Foundation, the OpenID Foundation, the Open Identity Exchange, and others including established SDOs such as W3C, IETF, ISO, etc. See [Appendix B](#) for more.
3. [Design Principles for the ToIP Stack](#) is the immediate predecessor to this specification (see the development tracks described in [Section 4](#)). It enumerates the set of design principles informing, guiding, and constraining the design of the ToIP stack. We especially recommend this document for a complete understanding of this specification.

As with all ToIP deliverables, the ToIP Foundation invites your feedback and suggestions. Please contact us via the [ToIP Foundation website](#).

## 2. Terminology

---

In this document, the key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL", when appearing in ALL CAPITALS, are to be interpreted as described in [IETF RFC 2119](#).

All other defined terms are linked to their definitions in the applicable ToIP [glossary](#) following the process defined by the [ToIP Concepts and Terminology Working Group](#)(CTWG). The glossaries cited in this specification include:

- [The ToIP Core Glossary](#)
- [The CTWG Glossary](#)
- [The TSWG Glossary](#)
- [The eSSIF-Lab Glossary](#)

*NOTE: Work on these glossaries is ongoing with the aim of being complete for the second public review draft (PR2).*

Terms that especially vital to this specification are also explained inline.

## 3. Motivations

---

*This section is informative.*

The goal of this specification is to define the overall requirements for a layered system architecture that enables interoperable trust relationships between any set of endpoints on the Internet. This is directly analogous to the role the TCP/IP stack plays in enabling interoperable data exchange between any set of endpoints on the Internet. The design patterns applicable to solving these interoperability challenges, and the motivations for each, are detailed at length in [Design Principles for the ToIP Stack](#).

Whether from the perspective of an implementer, a customer, or a policymaker, there are many benefits to a well-defined layered architecture:

- **Engineering stability.** The abstraction of bundling technologies and policies within distinct layers isolates changes within a layer from interactions and dependencies between layers. The result is a framework more resilient to structural changes than when the layers are not separated.
- **Wide adoption through a principled trust spanning layer.** The TCP/IP stack features a simple minimalistic protocol (the [Internet Protocol](#)) as a universal spanning layer. More diverse task-specific protocols are built on top of this spanning layer (e.g., TCP for connection-oriented, UDP for connectionless). In the case of ToIP, a trust spanning layer based on a well-grounded set of design principles can maximize adoption, interoperability and reachability.
- **Reliability.** A well-defined architecture enables the development of software components and applications that can be trusted to act in predictable, reliable ways—and that can expect other components and applications to do the same.
- **Interoperability and vendor independence.** As with the TCP/IP stack, the Bluetooth stack, the NFC stack, or other protocol stacks, implementations from multiple vendors can and should be interoperable, and customers should be able to switch between them while maintaining standardized functionality. In addition, whenever practical, ToIP should leverage existing technologies provided they are consistent with the [ToIP design principles](#).
- **Development communities.** A well-designed architecture stack helps spawn a robust, diverse community of developers building solutions whose interoperability depends on a core stack. More development attracts more innovation, more innovation attracts more adoption, producing a network effect benefiting the entire ecosystem.
- **Commoditization.** Standardization of a stack for mass adoption commoditizes it, reducing both the cost of implementation and the time to market. It also frees vendors to focus on their proprietary differentiation in their service offerings.

- **Public policy.** A well-defined architecture with clear and concise terminology helps policymakers and legal experts define coherent policies and regulations in a manner that serves the needs of society without constraining technical innovation and competition.

## 4. Audience, Purpose and Scope

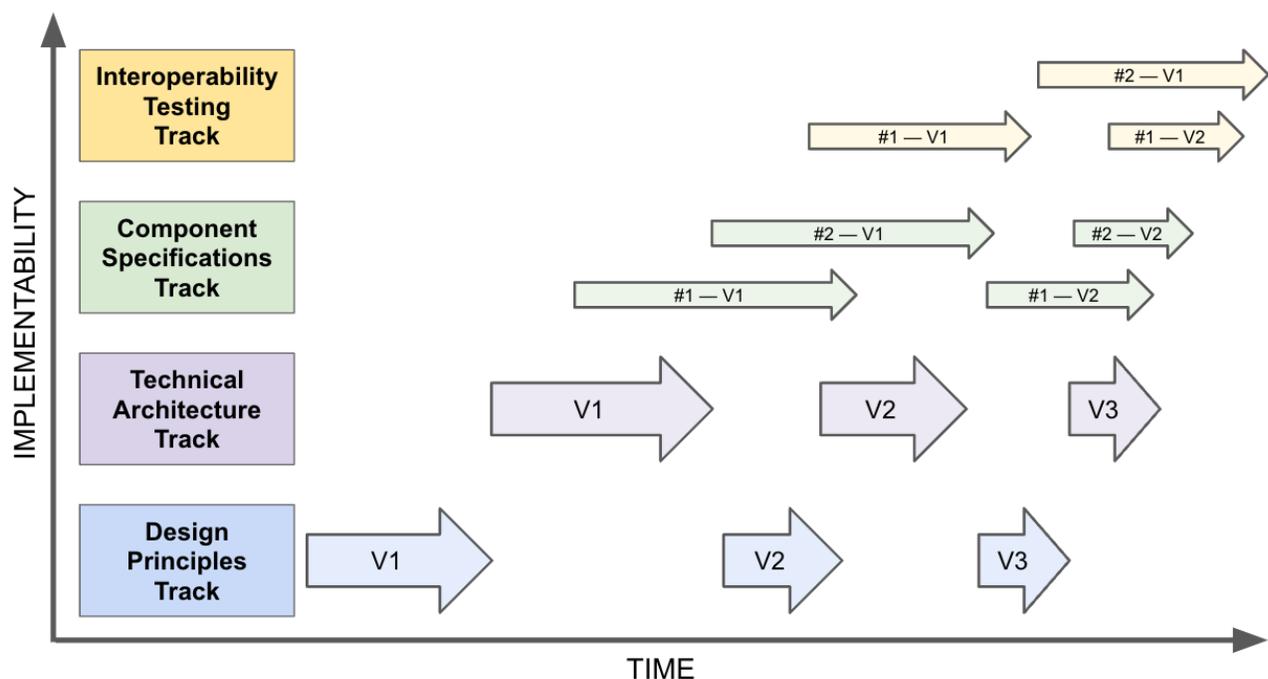
*This section is informative.*

The audience for this specification is protocol designers, system architects, software developers and product managers who wish to understand, influence, design, develop, or implement interoperable decentralized digital trust infrastructure, services, or applications.

The purpose of this specification is to define a reference architecture for the technology side of the ToIP stack, including the functions and behaviors required for each of the four layers and the functional and behavioral inter-dependencies between the layers:

- What each layer **must** do.
- What each layer *should* and *may* do.
- What behaviors are expected to support interoperability.
- What interactions each layer supports for other layers.

The goal of these architectural requirements is to inform subsequent development tracks as summarized in Figure 2:



## Figure 2: The planned progression of development tracks for the ToIP Technology Stack

For more information about the interrelationship and progression of these four development tracks, please see Appendix B: Evolution of the ToIP Stack.

The scope of this specification is limited to the Technical Architecture Track above, i.e., to defining the normative architectural requirements needed to guide the Component Specification Track. Success will be achieved if these requirements are sufficient to produce the component specifications needed to implement the architecture and prepare for the Interoperability Testing Track.

By focusing solely on the Technical Architecture Track, the following are explicitly out-of-scope:

1. The definition of specific protocols or interfaces at each layer (these will be produced in the Component Specifications Track).
2. The definition of specific interoperability profiles and test cases—including both vertical and horizontal interoperability—that can be used for commercial-grade test harnesses and testing labs (these will be produced in the Interoperability Testing Track).
3. The definition of specific intermediary systems or supporting systems for any layer.
4. The definition of specific applications (and their user interfaces) that run on top of the ToIP stack.
5. The definition of ToIP Governance Stack components such as trust frameworks or governance frameworks for usage of the ToIP stack within specific digital trust ecosystems.

*NOTE: We do not expect all of these additional deliverables, especially the component specifications, to be produced by the ToIP Technology Stack Working Group (or other ToIP Working Groups). Some have already been produced—and others are in development—by other standards development organizations (such as the Decentralized Identity Foundation, W3C, IETF, ETSI, and ISO), independent governing authorities, and independent developers.*

*NOTE: Due to the public policy implications, the ToIP Foundation is committed to producing a companion document called **ToIP Primer for Policymakers**. This document will guide policymakers, governing authorities, analysts, and other non-technical audiences who need to deeply understand the purpose, uses, and implications of the ToIP stack but do not need (or want) to dive into technical details.*

# 5. Example Use Cases

---

*This section is informative.*

*NOTE: Work on a set of example use cases illustrating the full breadth of usage of the ToIP stack is underway in the ToIP [Ecosystem Foundry Working Group](#). A summary table of these use cases is planned to be included in the second public review draft (PR2).*

# 6. Reference Architecture Overview

---

*This section is informative.*

## 6.1 Design Goals

---

A reference architecture of a complex system is an abstract framework consisting of a list of functional subsystems together with the interfaces and protocols needed to define the potential interactions and dependencies between these systems and/or external systems. This reference architecture provides a logical articulation of these interfaces and protocols which can then be translated into specific component specifications as described in Figure 2.

Such a reference architecture is an exercise in design guided by a set of most significant goals or principles. The overarching goals for the ToIP stack are twofold:

1. Define a general means of establishing trust between any two or more endpoint systems,
2. Achieve universal interoperability among implementations.

These twin objectives led the ToIP Foundation to begin the work with the Design Principles Track in Figure 2. In 2021, we developed a set of 17 [Design Principles for the ToIP Stack](#) that are the basis for the design choices reflected in this specification. For the full rationale behind each design principle, see [Design Principles for the ToIP Stack](#).

With regard to the first design goal, establishing trust between parties requires that each party develop confidence in the following properties of their relationship:

1. **Authenticity:** is the receiver of a communication able to verify that it originated from the sender and has not been tampered with?<sup>[1]</sup>
2. **Confidentiality:** is the contents of a communication protected so only authorized parties have access?

3. **Privacy:** will the expectations of each party with respect to usage of shared information be honored by the other parties?

Note that, in some trust relationships, confidentiality and privacy may be optional. Thus our design goal with the ToIP stack is to achieve these three properties in the order listed.<sup>[2]</sup>

With regard to the second design goal, the ToIP reference architecture shares the same goal of global scalability as the original Internet architecture. This involves several intertwined considerations that overlap and reinforce each other as summarized by the first four [Design Principles for the ToIP Stack](#):

1. The End-to-End Principle
2. Connectivity Is Its Own Reward (Universal Interoperability)
3. The Hourglass Model
4. Decentralization

## 6.2 The Four Layer Pattern

---

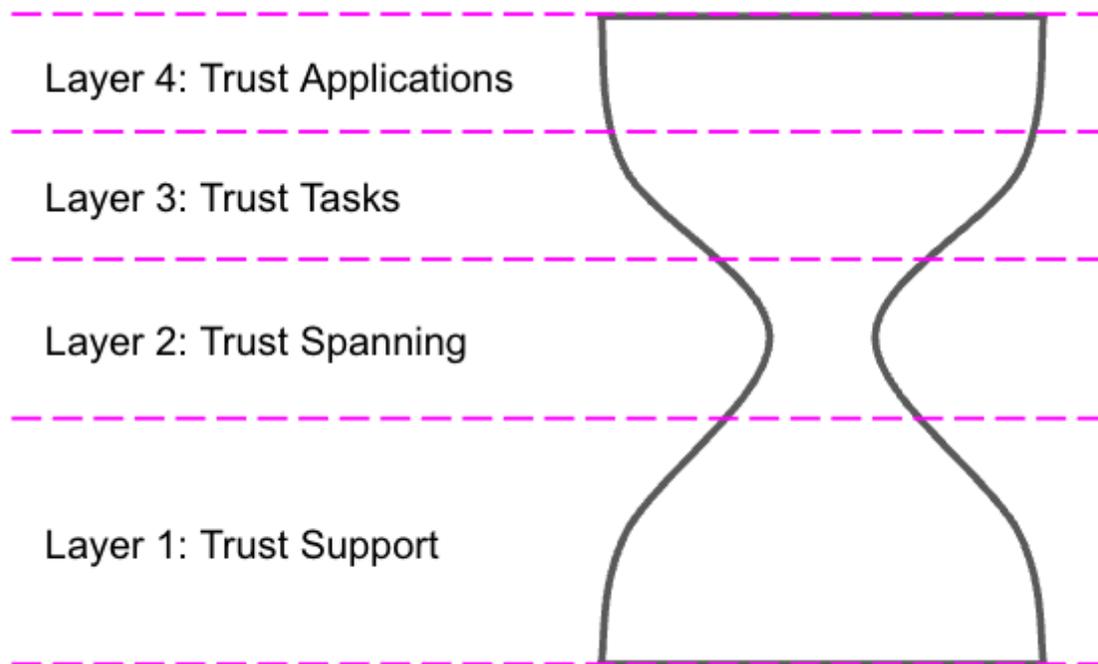
Together these considerations lead to the general *four-layer pattern* of a protocol stack summarized in Table 1.

Layer #	Generic Name	ToIP Name
4	Applications	Trust Applications
3	Supported protocols	Trust Tasks
2	Spanning protocol	Trust Spanning
1	Supporting protocols	Trust Support

**Table 1: The four layer pattern of large-scale protocol stacks**

The best-known example of this four-layer pattern is the TCP/IP Internet protocol stack, where any number of local area networking protocols at Layer 1 support a single spanning layer protocol—the IP protocol—at Layer 2. This spanning layer in turn supports multiple higher-level protocols at Layer 3 (e.g., TCP, UDP, HTTP, SMTP) designed to meet the needs of many different applications at Layer 4.

Much of the success of the Internet is attributed to this “hourglass” design in which the spanning layer protocol maximizes interoperability by providing a common way for all the higher level layers to communicate with all the lower level layers. This is why the design of the trust spanning layer should be “as simple as possible but no simpler”. Figure 3 illustrates how this same hourglass design applies to the four ToIP layers.



**Figure 3: How the four layer pattern fits the Hourglass Model**

For additional overviews of how the ToIP technology stack implements the hourglass model, see [Appendix B](#).

## 6.3 High-Level System Architecture

---

The reference architecture of the ToIP stack provides a generalization of different solutions to trust establishment over the Internet (or over other digital networks). This section introduces the basic concepts, requirements and vocabulary with which to consider: a) each functional component, b) the interface definitions and protocols between these components, and c) interoperability of solutions built upon those components. Subsequent sections will describe these components and protocols in more detail.

At the highest level, ToIP interactions occur between three basic types of interacting systems delineated by *locus of control*.

1. **Endpoint Systems** (often simply referred to as Endpoints): the ToIP systems between which end-to-end trust is enabled following the End-to-End Principle. See

## Section 7.1.

2. **Intermediary Systems** may be used to assist in the interactions between the Endpoint Systems. In that context, Intermediary systems are involved in the ToIP Trust Spanning Protocol, and may themselves be Endpoint Systems. Intermediary systems are not a dependency to Endpoint systems' trust relationship. See [Section 9](#).
3. **Supporting Systems** are typically required to support the definition of Endpoints and trust establishment between Endpoint Systems. Supporting Systems that facilitate the authenticity and autonomy of an Endpoint System are termed "Privileged" Supporting Systems, others are "Unprivileged" (see [Section 10.1](#)). Supporting Systems are not directly involved in the ToIP Trust Spanning Protocol. See [Section 10](#).

The relationships between these systems is shown in Figure 4.

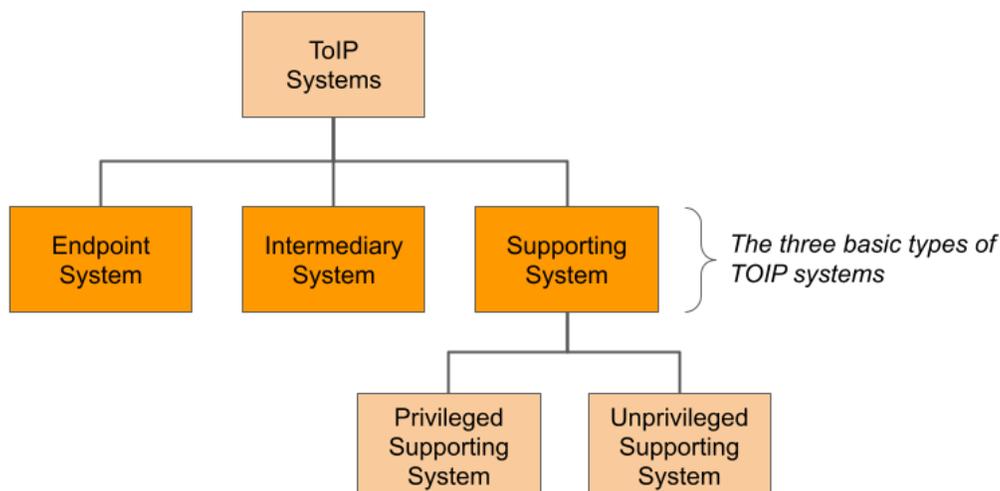


Figure 4: The three basic types of component systems in ToIP architecture

The definition of each system is anchored to its defined (and agreed) locus of control, i.e., who is able to exert control over the operation of that system. Clarity about the locus of control and the dependencies between systems is critical as end-to-end trust is constructed between any two Endpoint Systems. Each system, whether it is classified as an Endpoint System, Intermediary System or Supporting System, defines its own locus of control. An Endpoint System, for example, may be a tiny IoT device, a personal smartphone, or a large capacity service hosted in a cloud. The terms such as 'local or remote' or 'within, internal or outside' an Endpoint System should be understood as being with respect to its locus of control rather than physical location. What matters to the architecture is that it exhibits a consistent locus of control and, therefore, consistent interaction protocols with respect to other systems.

These subsystems collaborate with each other through three types of consistent ToIP interactions:

1. Endpoint System to Endpoint System
2. Endpoint System to Supporting Systems
3. Endpoint System to Intermediary Systems

ToIP Endpoint Systems and their interactions follow the 4-layer design pattern described in [Section 6.2](#). As we move up the stack (to Layers 3 and 4), the roles that may be played by an Endpoint System are often given more context-specific names. For example, at Layer 3, an Endpoint System involved in the Trust Task of exchanging [verifiable credentials](#) may be classified as an "Issuer", "Holder", or "Verifier" in that specific interaction context. These higher layer terms are specific to that context and must be consistent with the abstract general terms used in this reference architecture.

Figure 5 shows a high level view of how these three basic types of component systems might interact using the existing infrastructure of the Internet.

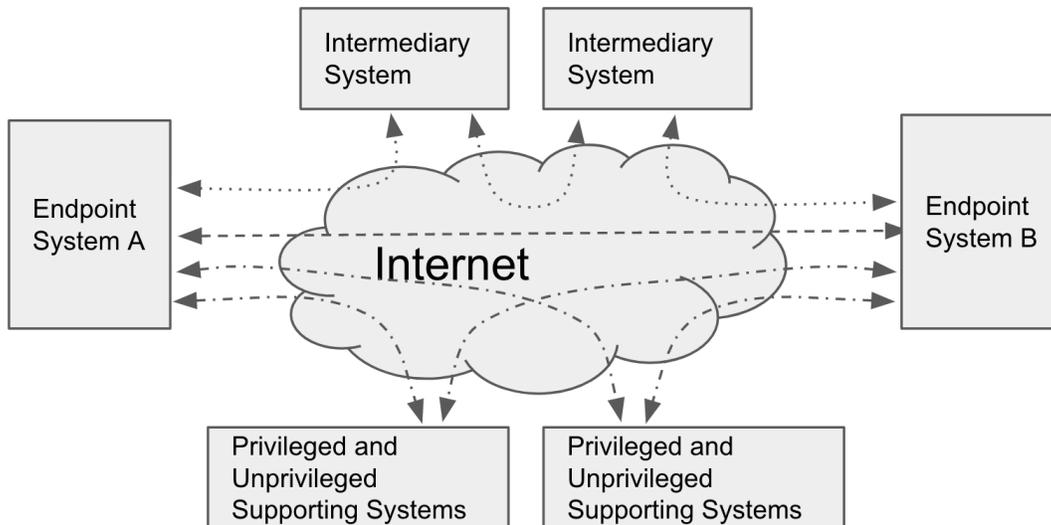
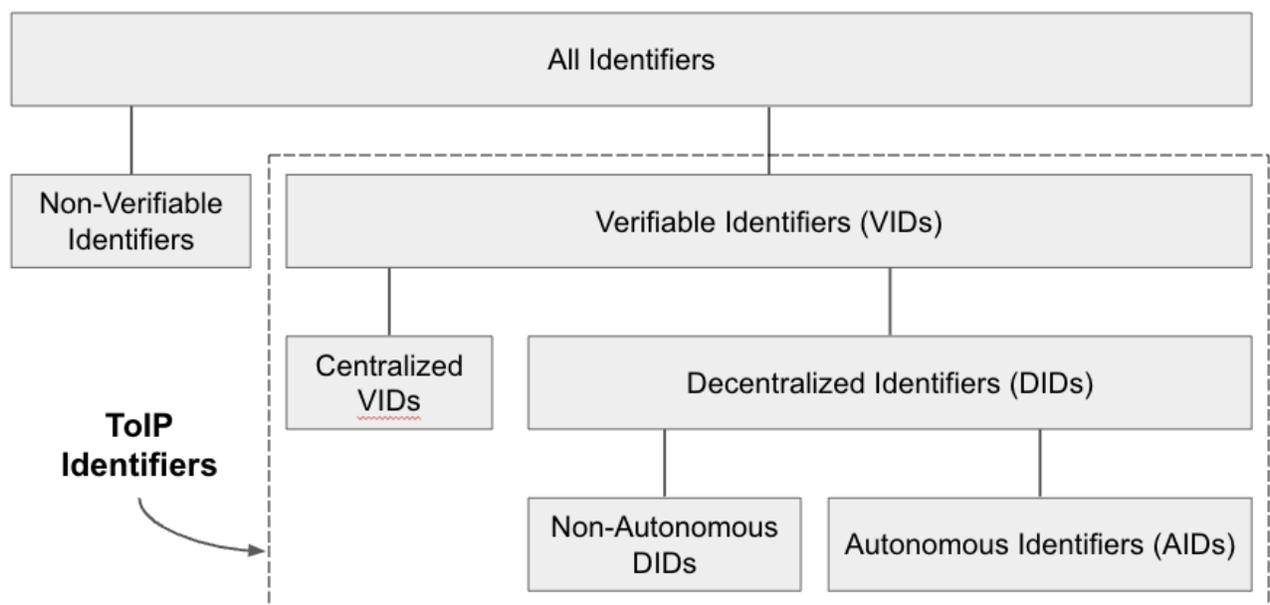


Figure 5: High level view of ToIP consistent system interactions

The normative requirements for each type of subsystem and interaction across the ToIP layers are specified in the following sections.

## 6.4 ToIP Identifiers

Just as IP addresses are the heart of the Internet TCP/IP stack, ToIP identifiers are the heart of the ToIP stack. Figure 6 illustrates the three basic types of ToIP identifiers.



## Figure 6: The three classes of identifiers used with ToIP architecture

Design Principle #5 ([Cryptographic Verifiability](#)) states that “messages and data structures exchanged between parties should be verifiable as authentic using standard cryptographic algorithms and protocols”. This requires that Endpoint Systems be able to associate, discover and verify the cryptographic keys associated with a ToIP identifier. This specification will refer to identifiers that meet this basic requirement of cryptographic verifiability as verifiable identifiers (VIDs).

VIDs divide into two subclasses: those that rely on centralized registries, and those that do not. An example of the former is the authority portion of an HTTPS URL. It relies on a centralized DNS registry for resolution, but that domain name is associated with an X.509 digital certificate that provides a cryptographic binding with a public key.

VIDs that do not require a centralized registry are known as decentralized identifiers or DIDs. DIDs fulfill the requirement of Design Principle #4 ([Decentralization by Design and Default](#)). The W3C has established a global standard for the syntax and resolution format of [DIDs](#).

DIDs can be further subdivided into two classes: those whose verification requires the use of an external system outside of the autonomy boundary of the DID controller, and those whose verification can be accomplished using cryptography alone. The latter are often called self-certifying identifiers. Since self-certifying identifiers enable an Endpoint System to be completely autonomous, this specification will refer to this class of identifiers as autonomous identifiers (AIDs). A [KERI autonomic identifier](#) is an example of an AID.

Requirements for ToIP identifiers are covered in [Section 8.2](#).

# 7. Endpoint Systems and the Layered Stack

---

*This section is normative.*

## 7.1 Endpoint Systems

---

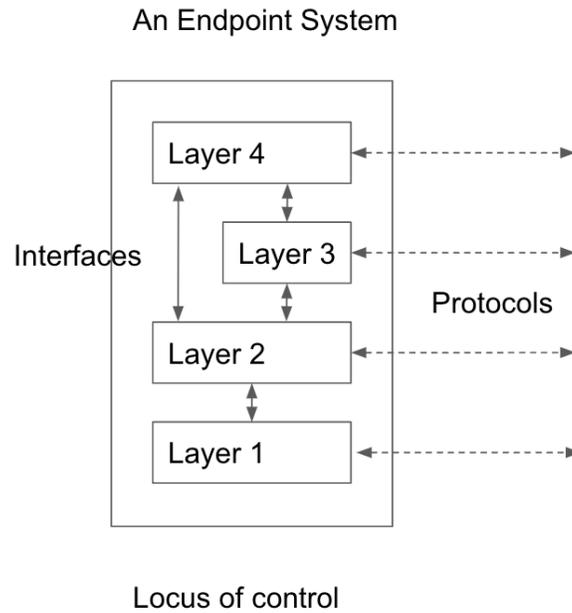
Endpoint Systems represent ToIP systems that are under a party's direct control. An Endpoint System's boundary is delineated by its locus of control. A party means the entity that is evaluating, relying on, and benefiting from a trust relationship. In other words, a party is any user of the system without regard to their role in the system. This represents a contrast with the traditional identity and access management (IAM) distinct roles of a user who is making trust assertions and a *relying party* who is relying on those assertions to make a trust decision. In a ToIP system, Endpoint Systems have a symmetric *peer-to-peer* trust relationship in Layer 2 — the trust spanning layer.

Endpoint Systems are autonomous in the sense that a party's locus of control is the whole Endpoint System by definition. This means a potential compromise of other Endpoint Systems, Intermediary Systems, or Supporting Systems will not directly compromise the integrity of a given Endpoint System. Each Endpoint System can be simple or very complex, i.e., it may have many further divided functions and/or services, however in this reference architecture, we shall consider the abstract Endpoint System autonomous. Implementers SHOULD ensure autonomy for Endpoint Systems [REQ A.1]

Common examples of Endpoint Systems include:

- A mobile phone owned and administered by an individual.
- A merchant's web server (on-premise or in the cloud), administered by the merchant.
- A financial institution's distributed digital services, including certain online services for trust functions, administered by the financial institution.
- An IoT pollution sensor device owned and administered by a city.

Befitting Design Principle #1 ([The End-to-End Principle](#)), Endpoint Systems are the ultimate targets of the requirements of ToIP architecture. They are likely to be much larger in number — by several magnitudes — compared to Intermediary or Supporting Systems. They implement most of the functions in the ToIP architecture and represent the biggest challenge for interoperability and scalability.



**Figure 7: Endpoint System**

Within an Endpoint System's locus of control, a higher layer uses the functions of a lower layer through an **Interface**. In ToIP architecture, functions within an Endpoint System are decomposed into layers in a vertical stack where layer boundaries are defined by their corresponding Interfaces. In a ToIP Endpoint System, the higher layers of the ToIP protocol stack **MUST** communicate with the lower layers via defined interfaces. [REQ A.2]

In addition to the internal layer interfaces implemented by hardware and software resources within the Endpoint System's locus of control, an Endpoint System may also rely on the services of other Supporting Systems that are located outside of the Endpoint System's locus of control but accessible through the Internet to perform their functions. This type of interaction requires a defined **Protocol**.

The distinction between an Interface and a Protocol is whether the systems communicating over the protocol represent different loci of control. For example, simply distributing the functions within a particular layer over the Internet — such as having some of the functions performed using cloud computing or web services—does not necessarily require a defined protocol if all of the functions are under the same locus of control. However an agreed protocol may be necessary if the communicating systems are under different loci of control. What is essential is delineating who has control over what in order to reason about trust relationships.

The four layer stack within an Endpoint System is defined in the following sections.

## 7.2 Layer 1: Trust Support

---

If a ToIP Endpoint System includes Trust Support Functions *within its locus of control*, then those functions **MUST** be included at Layer 1 of the Endpoint System. [REQ L1.1] The exact nature of the Trust Support Functions required by any particular Endpoint System may vary significantly depending on the Endpoint System's physical manifestation and numerous other design goals (e.g. cost, location, convenience, power usage, reliability and so on). For example the Trust Support Functions required for a full-featured smartphone vs. a cloud server vs. an IoT thermostat may be very different.

Examples of Trust Support Functions designed to specifically support machine-to-machine trust (aka cryptographic trust or technical trust):

- Cryptographic hardware modules capable of generating good quality cryptographic key materials.
- Sufficiently secure storage of secrets and cryptographic materials.
- Sufficiently secure computing environment.
- Sufficient communication functions for the intended deployment environment.

*NOTE: while this specification generally assumes the Internet as the common networking environment, Internet support is not strictly required. The ToIP stack may be implemented over any communication medium capable of supporting the communication functions.*

Examples of Trust Support Functions designed to specifically support human-to-human trust (aka business trust or legal trust) include:

- Identity binding mechanisms that associate a natural person to the Endpoint System itself, or to data artifacts on the Endpoint System such as identity claims or verifiable credentials. One of the strongest identity binding mechanisms is biometrics — physiological or behavioral characteristics that identify the individual (facial recognition, fingerprint readers, voice recognition, palm recognition, and so on). Layer 1 would provide the hardware and software support for registering, storing, and processing biometric primitives.
- Hardware-based trust attestation systems, such as those incorporated within [Trusted Platform Modules](#), and other confidential computing systems that can provide legally valid evidence of the security characteristics of a Layer 1 component.

Diversity of implementations of Layer 1 Trust Support Functions is *intentional* and a key goal of the ToIP stack design.

*NOTE: For functional, performance, security, or other reasons, a Layer 1 Trust Support Function implementation may use a remote service outside its locus of control, e.g., a distributed ledger, distributed directory, distributed database, distributed file system, or distributed hash table. These systems are Supporting Systems to the Layer 1 implementation; they are not part of Layer 1 itself. See [Section 10.1](#).*

## 7.3 Layer 2: Trust Spanning

---

Layer 2 is the trust spanning layer of the ToIP stack. In keeping with Design Principle #3 ([The Hourglass Model](#)), this means there is only one requirement for Layer 2: A ToIP Endpoint System **MUST** communicate with another ToIP Endpoint System using the ToIP Trust Spanning Protocol. [REQ L2.1] No other functions are required.

The requirements for the ToIP Trust Spanning Protocol are defined in [Section 8](#).

## 7.4 Layer 3: Trust Tasks

---

Many applications may require more complex trust-building functions than the minimal set offered directly by the ToIP Trust Spanning Protocol. When one of these functions is reusable across multiple contexts that are separated in time, space, or perspective, we call it a Trust Task. Trust Tasks can be standardized as their own higher-level protocols at Layer 3 of the ToIP stack.

A Layer 3 Trust Task Protocol **MUST** communicate either over the Layer 2 ToIP Trust Spanning Protocol or over another Layer 3 Trust Task Protocol for all communications related to trust establishment between Endpoint Systems. [REQ L3.1] This is directly analogous to how TCP and UDP communicate over IP, and how HTTP communicates over TCP. A Layer 3 Trust Task **MAY** use other protocols, but only for other purposes (since short-circuiting Layer 2 when establishing trust with other Endpoint Systems would undermine the trust guarantees of the ToIP stack). [REQ L3.2]

Note that because Confidentiality and Privacy are optional for the Layer 2 ToIP Trust Spanning Protocol, the following requirement applies: A Layer 3 Trust Task Protocol intended to communicate private data **SHOULD** support Confidentiality and Privacy. [REQ L3.3]

There can be as many Trust Tasks Protocols as are needed by Layer 4 Trust Applications. Some examples of Trust Tasks include:

- Human authentication (as opposed to cryptographic authentication performed at Layer 2), including biometric authentication

- Exchanges of verifiable credentials (issue, request, offer, present, revoke)
- Provisioning, updating, and verification of digital identities
- Consent management
- Requesting and signing of digital documents
- Secure messaging
- Secure data sharing
- Digital payment or value exchange in any form
- Digital auctions
- Digital notaries

## 7.5 Layer 4: Trust Applications

---

Layer 4 is an open-ended application layer for any application that needs to engage in trusted interactions. Layer 4 Trust Applications MAY use any number of Layer 3 Trust Task Protocols. [REQ L4.1].

If a Layer 4 Trust Application does not use a Layer 3 Trust Task Protocol, it MUST communicate with other Endpoint Systems using the Layer 2 Trust Spanning Protocol. [REQ L4.2]

Layer 4 is the layer where humans “touch” the ToIP stack, so this is where Design Principle #8 ([Trust is Human](#)) and #14 ([Trust and Technology have a Reciprocal Relationship](#)) come into play. The human experience of digital trust is so critical that Layer 4 has one more requirement: A Layer 4 Trust Application MUST support any ToIP-defined Trust Affordances relevant to that application. [REQ 4.3]

# 8. The ToIP Trust Spanning Protocol

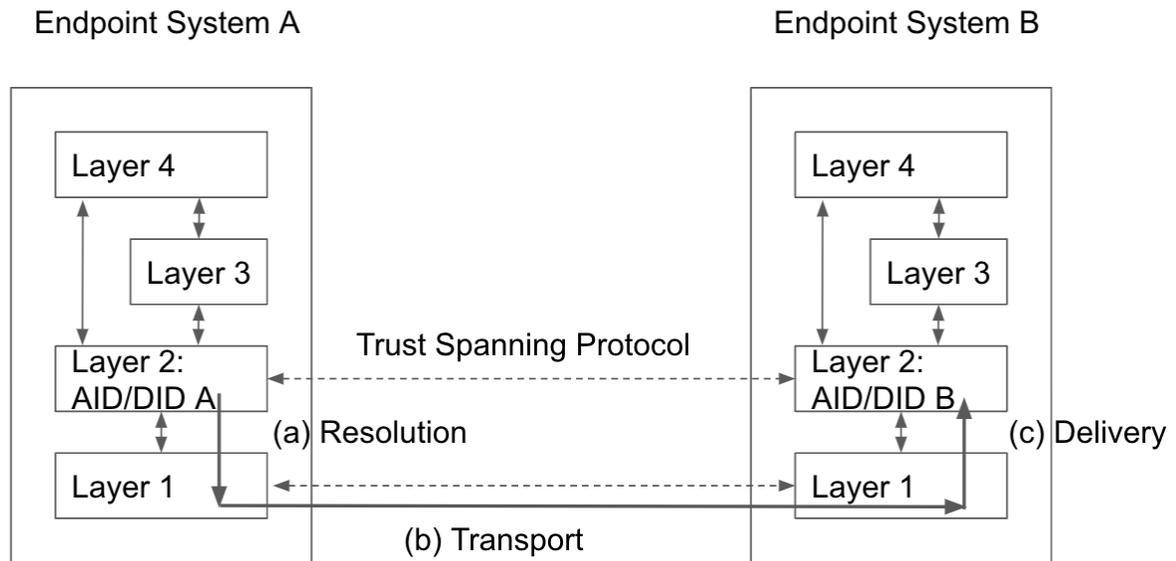
---

*This section is normative.*

## 8.1 Overview

---

This section describes the ToIP Trust Spanning Protocol required at Layer 2 to communicate between any two Endpoint Systems. The overall protocol operation is shown in Figure 8 below.



**Figure 8: Overview of the ToIP Spanning Protocol**

The main function of this protocol is to enable universal end-to-end communication among all Endpoint Systems using trusted messages. This architectural choice is based on the following considerations:

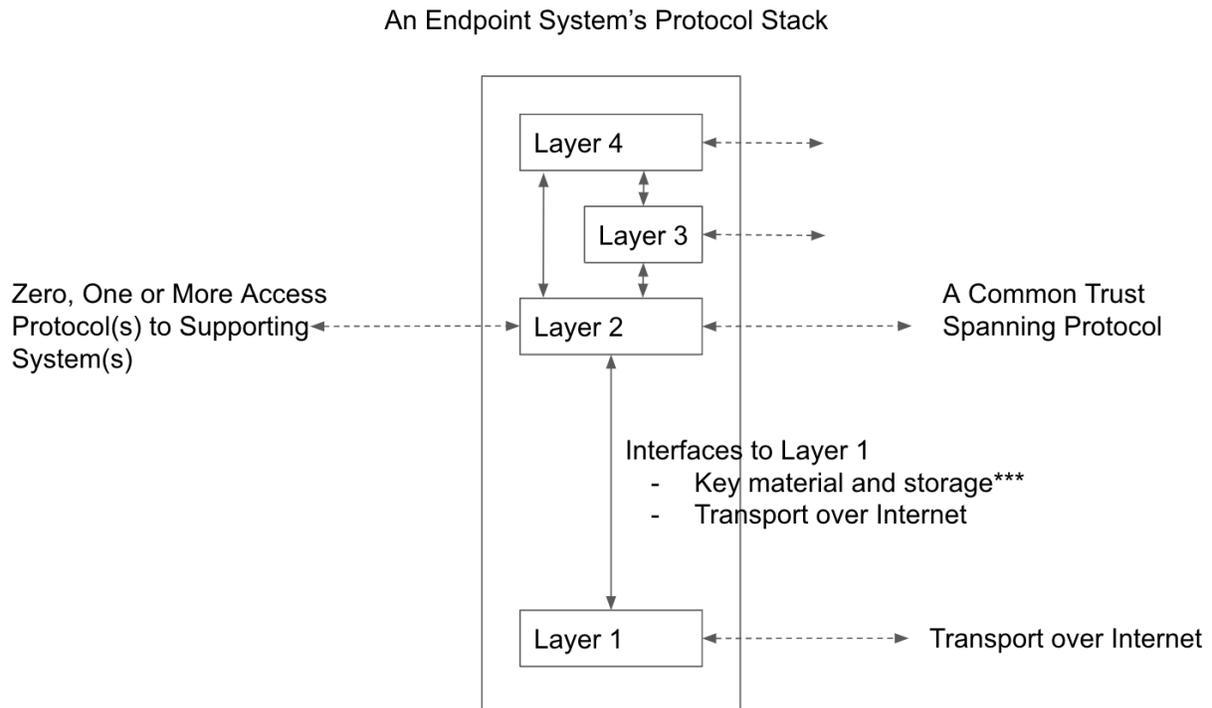
- Existing Internet, local network, and mesh network infrastructure already supports universal end-to-end communication through various types of transport mechanisms.
- This form of communication should be able to be implemented on all common types of Endpoint Systems with minimum overhead and constraints.
- Messaging-based communication is a least common denominator that provides sufficient foundation for building more complex trust functions at higher layers.

This protocol is designed to be universal in the sense that all Endpoint Systems, regardless of their form factors or implementation methods, can communicate with each other using messages incorporating standard trust guarantees.

To achieve ubiquity, this protocol should be kept as simple as possible to ease implementation challenges and allow maximum flexibility on all variants of Endpoint Systems. Thus the requirements in the following sections are not only necessary but sufficient. Strong preference must be given not to add additional functions to this protocol unless they are universally beneficial. Strong preference must also be given to a single common protocol specification for maximum any-to-any interoperability.

A view of the ToIP protocol stack on an Endpoint System is shown in Figure 9. The component specification for the ToIP Trust Spanning Protocol therefore needs to specify:

1. How to generate and maintain identifiers with the properties described in [Section 6.4](#).
2. The common message format that meets the design goals described in [Section 6.1](#).
3. How lower layer transport protocol(s) can be used to deliver messages between Endpoint Systems.
4. Any required support from ToIP Layer 1.



**Figure 9: A view of the ToIP protocol stack on an Endpoint System**

The following sections enumerate the requirements in each of these four areas.

## 8.2 Identifiers

A key difference between Internet architecture and ToIP architecture is that the former only needed to identify the network endpoints of devices for data communications. The solution was Internet Protocol (IP) addresses: a global addressing scheme for network endpoints independent of any specific local area network.

By contrast, ToIP Layer 2 architecture needs to identify and route messages between the entities participating in trust relationships. While this set of entities may include the devices serving as Endpoint Systems, it extends beyond the network to include identifiers for parties —people and organizations using the network to interact and transact.

In order to establish trust in ToIP identifiers ([Section 6.4](#)) regardless of the type of entity to which they are bound, they must meet the following requirements:

- A ToIP identifier **MUST** be unique within the context in which it is used for identification. [REQ L2.2] This means ToIP identifiers intended to be globally resolvable need to be globally unique; identifiers intended to be locally resolvable need to be locally unique.
- A ToIP identifier **MUST** be a verifiable identifier, i.e., verifiably bound to at least one set of cryptographic keys discoverable via an associated discovery protocol. [REQ L2.3]
- A ToIP identifier **SHOULD** be a decentralized identifier, i.e., a verifiable identifier that does not require registration with a centralized authority. [REQ L2.4]
- A ToIP identifier **SHOULD** be an autonomous identifier, i.e., a decentralized identifier that is self-certifying and fully portable. [REQ L2.5]
- A ToIP identifier **SHOULD** support rotation of the associated cryptographic keys for the lifetime of the identifier. [REQ L2.6]
- A ToIP identifier **MAY** also support rotation to an entirely different ToIP identifier that can be cryptographically verified to be a synonym of the original ToIP identifier. [REQ 2.7]
- A ToIP identifier **SHOULD** support the ability to: a) associate the identifier with the network address of one or more ToIP Systems that can deliver to one or more Endpoint Systems under the locus of control of the ToIP identifier controller, and b) if desired by that controller, enable that association to be discoverable. [REC L2.8]

Special considerations apply when a ToIP identifier needs to be provably bound to a specific party, i.e., a person or an organization. Proof of such a binding can be a critical factor in establishing a desired level of assurance in the identity of that party. Such proof can be accomplished using multiple mechanisms such as:

1. Proof of control of the cryptographic keys bound to the ToIP identifier.
2. Proof of control of one or more verifiable credentials describing the identified party.
3. Proof of one or more biometric primitives describing the identified party.

Such proofs may require support from one or more Layer 1 Trust Support Functions within the Endpoint, and/or support of one or more Supporting Systems outside of the Endpoint, and/or the additional invocation of one or more Layer 3 Trust Task Protocols. These steps are out-of-scope for the Layer 2 Trust Spanning Protocol.

Different considerations apply when a ToIP identifier needs to be provably bound to a digital resource, such as a file, photo, or video. This can be accomplished using ToIP identifiers that serve as content-addressable identifiers or self-addressing identifiers that are derived from a cryptographic hash of the subject resource. For example, see [ACDC](#).

## 8.3 Messages

---

Messages are the lingua franca of the ToIP Trust Spanning Protocol. To achieve the design goals in [Section 6.1](#), the following requirements must be met:

The ToIP Trust Spanning Protocol specification **MUST** define how to construct and format messages that are cryptographically verifiable to have the following four properties:

- **Authenticity:** the message was sent from a sender who has control over the ToIP identifier.
- **Integrity:** the contents of the message transmitted by the sender are received by the recipient without modification.
- **Confidentiality:** the contents of the message are only accessible by authorized parties.
- **Privacy:** the contents of the message are bound to conditions of usage agreed to by the parties. [REC L2.9]

In a ToIP Endpoint System, an implementation of the ToIP Trust Spanning Protocol **MUST** support authenticity and integrity. [REQ L2.10]

In a ToIP Endpoint System, an implementation of the ToIP Trust Spanning Protocol **MAY** support confidentiality and privacy. [REQ L2.11]

The ToIP Trust Spanning Protocol **MUST** enable the composition of higher-level Trust Tasks (such features as co-protocols). [REQ 2.12] Examples of such features include discovery, threading, timeouts, ACKs, and attachments. However this must be balanced with the requirement to only add additional functions to this protocol if they are universally beneficial.

The ToIP Trust Spanning Protocol **MUST** support extensible message schema. [REQ 2.13] This enables different Trust Task protocols to be constructed without changing the base format.

## 8.4 Routing

---

Routing of a message from a sender to a receiver proceeds in three steps as shown in Figure 8:

1. **Address resolution** takes the ToIP identifier of the receiver and resolves it to: a) the network address of an Endpoint System for the receiver that supports the desired Layer 1 transport mechanism, and b) the associated cryptographic keys. For example, if the ToIP identifier is a DID and the desired transport is HTTP, then a DID resolver resolves the DID following the associated DID method to retrieve the DID document. It then selects: a) the service type associated with the ToIP Trust Spanning Protocol and extracts an HTTP URL to which a connection can be made to deliver the message to the other Endpoint System, and b) the required cryptographic keys.
2. **Transport** is the Layer 1 mechanism to send the message to the Endpoint System of the receiver or to an Intermediary System which can eventually deliver the message. In the above example, HTTP is the transport. Over the Internet, any transport layer protocol may be a suitable transport. Other contexts may use other transports, e.g. Bluetooth, QR code, or a publish-subscribe messaging system.
3. **Delivery** is the final step of delivering the message to Layer 2 of the receiver's Endpoint System. This step may include a sub-step for an Intermediary System ([Section 9](#)) to deliver the message to the Endpoint System, and a second sub-step for the Endpoint System's Layer 1 transport to deliver the message to the Layer 2 interface.

These steps lead to the following requirements:

The ToIP Trust Spanning Protocol MUST support resolution of ToIP identifiers to: a) the network addresses of receiving Endpoint Systems, and b) any required cryptographic keys. [REC 2.14]

The ToIP Trust Spanning Protocol MUST support transport of messages via ToIP Layer 1 interfaces. [REC 2.15]

The ToIP Trust Spanning Protocol MUST support delivery of messages to the Layer 2 interface of the Endpoint System of the ultimate receiver of the message. [REC 2.16]

The ToIP Trust Spanning Protocol MUST support the option to deliver messages via Intermediary Systems. [REC 2.17]

The ToIP Trust Spanning Protocol MUST support confidentiality with regard to the metadata required for message routing. [REC 2.18]

## 8.5 Interface to Layer 1

---

Given these requirements for the ToIP Trust Spanning Protocol at Layer 2, the Trust Support Function interfaces at Layer 1 should only need to include the following. Note that Layer 3 Trust Tasks or Layer 4 Trust Applications may also need to call these interfaces directly.

1. **Key Management System (KMS)** is the interface for generating cryptographic quality keys, random numbers, or other values required by the cryptographic primitives used by the protocol.
2. **Secure storage** is the interface through which Layer 2 can create, read, write, and delete confidential or secret data.
3. **Transport** consists of one primitive via which the sender's Layer 2 implementation can submit a message for transmission and another primitive through which the receiver's Layer 1 implementation can deliver a message up to Layer 2.
4. **User binding** is the interface via which a Layer 2 implementation can request and verify a biometric or other authentication information from a user.

## 9. Intermediary Systems

---

*This section is normative.*

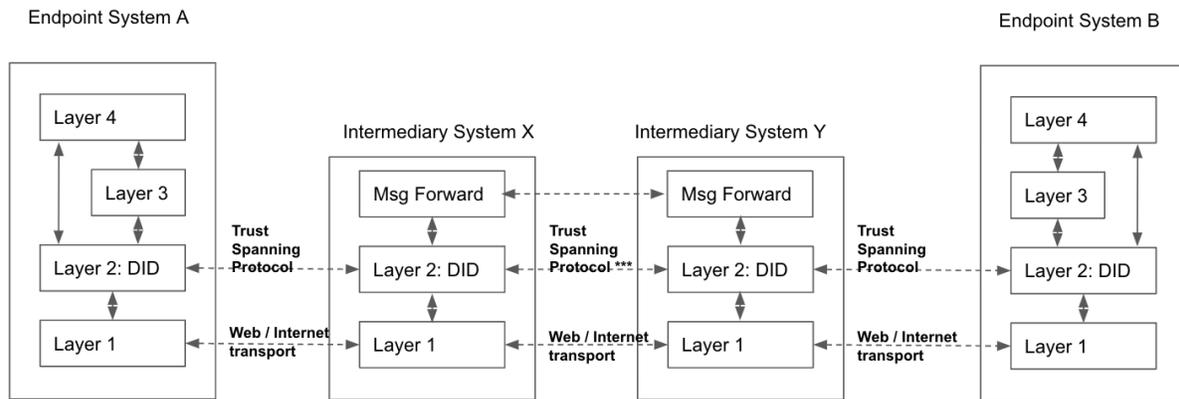
Intermediary Systems are mediators for facilitating the ToIP Trust Spanning Protocol. Since the Internet itself is routable as long as a ToIP identifier can be resolved to a unique IP address, Intermediary Systems are not absolutely required. However they can be very beneficial in other aspects.

Examples of useful Intermediary Systems include:

- **Store and forward intermediaries.** Some Endpoint Systems are not always connected to the Internet (e.g. smartphones). As with email, effective communication between an Endpoint System and a smartphone or similar device could use an Intermediary System to receive and store the messages while the device is not connected. Messages are subsequently delivered to the eventual receiver when the device becomes connected again.
- **Multi-device intermediaries.** Individuals who use multiple computing devices (e.g., smartphone, laptop, smart watch, smart car, etc.) may choose to use an intermediary to simplify configuration and management of these devices as well as routing of messages to the appropriate device.
- **Anonymizing intermediaries.** When confidentiality is desired in the routing of messages, intermediaries designed for this purpose can obfuscate the message routing path.

- **Auditing intermediaries.** When auditing of message traffic is required for regulatory or compliance reasons, intermediaries designed for this purpose can maintain the necessary audit logs.

Intermediary Systems differ from Supporting Systems because they reside between Endpoint Systems and are visible to the Endpoint Systems.



**Figure 10: The role of Intermediary Systems**

In Figure 10, end-to-end communication between Endpoint Systems A and B are routed through Intermediary Systems X and Y. In this case, all systems implement the Layer 2 protocol as described in [Section 8](#). Routing uses “nested envelopes” as follows:

1. Endpoint System A prepares a message for Endpoint System B and puts it in an inner message envelope addressed to Endpoint System B.
2. Endpoint System A places the inner message envelope inside an outer message envelope addressed to Intermediary System X.
3. Endpoint System A delivers the outer message envelope to Intermediary System X.
4. Intermediary System X removes the outer message envelope and replaces it with a new outer message envelope addressed to the next hop: Intermediary System Y.
5. Intermediary System X delivers the new outer message envelope to Intermediary System Y.
6. Intermediary System Y removes the outer message envelope.
7. Intermediary System Y delivers the inner message envelope to Endpoint System B.

This pattern casts one requirement for the use of Intermediary Systems:

A ToIP Intermediary System SHOULD be able to perform the functions of a ToIP Endpoint System for the purpose of routing enveloped messages using the ToIP Trust Spanning Protocol. [REC A.3]

# 10. Supporting Systems

---

*This section is normative.*

## 10.1 Overview

---

An Endpoint System may utilize services from any number of Supporting Systems, either privileged or unprivileged, over the Internet or other networks.

- **Privileged Supporting Systems** are integral to the dependent Endpoint System's autonomy and authenticity. A privileged Supporting System must implement strongly qualified trust mechanisms in order to play this role. Such trust mechanisms can be a combination of technology (e.g. algorithmic) and governance policies (see [Section 12](#)). For example, a blockchain is a privileged supporting system for a DID method whose root of trust is the blockchain.
- **Unprivileged Supporting Systems** are Supporting Systems that are not required to support an Endpoint System's autonomy and authenticity. For example, a website that serves as a non-exclusive convenient discovery (e.g. advertising or search) mechanism for public DIDs is unprivileged.

Each type of Supporting System may have a service access protocol standardized for the type of service it offers. There may be many such services with many different protocols. One Endpoint System may utilize one set of Supporting Systems while another Endpoint System may use a different set of Supporting Systems. This difference in the types of Supporting Systems used does not impede the two Endpoint Systems in interoperating through the Layer 2 ToIP Trust Spanning Protocol. Therefore, standardization across different services is not required.

An example of a common protocol stack for this purpose is a defined Web Service running on top of HTTPS. However, many types of protocols may be used for different Supporting Systems.

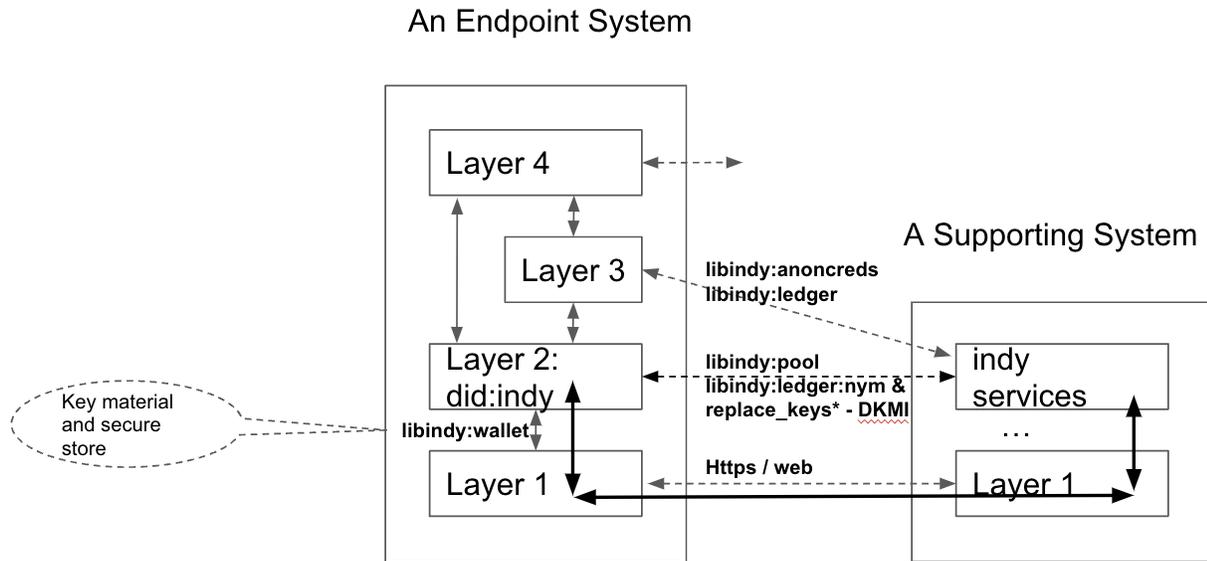
The ToIP protocol stack in an Endpoint System MAY use the services of a Supporting System at any layer. [REC A.4] Such design decisions can be made layer by layer to optimize the functions performed in each layer.

The following sections illustrated the layered interaction between Endpoint Systems and Supporting Systems using examples of known implementations.

## 10.2 Example 1 - A DID Method

---

A DID Method may be implemented based on a distributed ledger, e.g. Hyperledger Indy. An Endpoint System, in this example, may be implemented using a Hyperledger Aries agent software module running on either a mobile device or a cloud platform. The Indy ledger is a Privileged Supporting System and the Aries agent implements layer 2 and layer 3 of the Endpoint System stack. Such a design pattern is illustrated in Figure 11.



**Figure 11: An example of Hyperledger Indy as a Supporting System**

A Layer 2 implementation must implement both DID resolution and the ToIP Trust Spanning Protocol. To implement DID resolution in this example, the Aries agent uses a local service (i.e. within its locus of control), i.e. a digital wallet, which relies on, eventually, a KMS function and a secure storage function within the Endpoint System. It also uses a remote service (i.e. outside of its locus of control) — the Indy blockchain — via web service APIs built on top of HTTPS and other web protocols. This remote service protocol consists of three components in the case of Aries-Indy: pool API, anoncred API, and payment API. The web service eventually relies on the Internet Protocol stack for routing, transport and delivery. Collectively, it is a complete Endpoint System to Supporting System Protocol that in this case runs over the web.

## 10.3 Example 2 - A KERI Witness



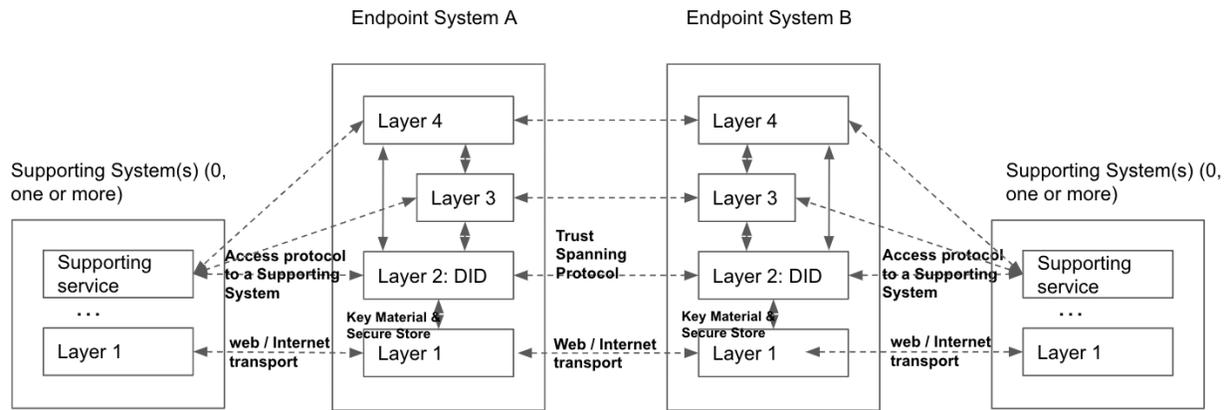


Figure 13: A generalization of how Endpoint Systems and Supporting Systems interact

## 11. Endpoint System Interoperability

### 11.1 Interoperability between Endpoint Systems Using Decentralized Identifiers

Section 6.4 states that “Endpoint Systems [need to] be able to associate, discover and verify the cryptographic keys associated with a ToIP identifier”. This capability is essential in order for two or more Endpoint Systems to be able to discover and connect with each other over the ToIP Trust Spanning Protocol.

If an Endpoint System is identified with a publicly resolvable decentralized identifier (DID) as defined in section 6.4, this is straightforward because a DID resolver can:

1. Resolve the DID to the authoritative DID document.
2. Extract the appropriate public key.
3. Extract the service endpoint URI for the ToIP Trust Spanning Protocol.

If an Endpoint System is identified with a private, pairwise DID — called a **peer DID** — the discovery and exchange of a DID document needs to use an out-of-band interaction (OOBI) protocol. Common examples include QR codes and custom-generated deep links.

### 11.2 Interoperability between Endpoints Using Other Verifiable Identifiers

If an Endpoint System is not identified with a DID, but with some other kind of verifiable identifier (VID) as defined in [section 6.4](#), then a different approach must be used to bootstrap communications using the ToIP Trust Spanning Protocol. This requires enabling discovery and verification of:

1. The authoritative public key for the Endpoint System.
2. The authoritative service endpoint URI for communicating with the Endpoint System over the ToIP Trust Spanning Protocol.

If the VID is an HTTPS URL, there are at least two solutions:

1. Conversion of the HTTPS URL into a **did:web:** identifier as described in the [ToIP X.509 PKD Interop page](#).
2. Issuance by a trusted issuer (such as a certification authority) of a verifiable credential whose subject is the HTTPS URL and whose claims assert the authoritative public key and ToIP Trust Spanning Protocol service endpoint URI.

We anticipate that integration of decentralized PKI and X.509 PKI will be a topic of increasing interest and innovation.

## 12. Integration with the ToIP Governance Stack

---

As explained in the Introduction, this specification, maintained by the ToIP [Technology Stack Working Group](#), is focused entirely on requirements for the ToIP Technology Stack. A separate set of specifications, maintained by the ToIP [Governance Stack Working Group](#), defines the requirements for the ToIP Governance Stack. The first generation of the ToIP Governance Architecture Specifications were published in January 2022 and are summarized [here](#).

Although the ToIP Governance Architecture Specifications consist largely of recommendations about the structure and content of governance documents for ToIP-based digital trust ecosystems, there are a very small but vital set of technical requirements that are essential for “tying the two stacks together”.

In particular, section 3 of the [ToIP Governance Architecture Specification V1.0](#) specifies a set of **identification requirements** for ToIP-compatible governance frameworks. A high-level summary:

1. The **primary document** for the governance framework **MUST** be assigned a DID and be retrievable via a DID URL. This DID identifies the governance framework itself as

- a digital object, and the DID URL allows it to be viewed and verified by any party.
2. All other **controlled documents** in the governance framework **MUST** have DID URLs.
  3. The DID URLs for all governance framework documents **MUST** be **versioned** as the documents are versioned.
  4. The **governing authority, administering authority** (if separate from the governing authority), and all **governed parties** in the governance framework **MUST** be identified with DIDs.

The use of persistent, discoverable, cryptographically verifiable identifiers for all parties and documents governing a digital trust ecosystem makes it much easier to bind technology to policy within a digital trust ecosystem. For example:

- A verifiable credential issued within the ecosystem can include a claim asserting the DID of the authoritative governance framework.
- The credential can also include the DID of one or more trust registries where a holder or verifier can verify that the DID of the credential issuer is authorized to issue that particular type of credential under that governance framework.
- A credential asserting certification under the governance framework can include the DID of the governance framework, the DID of the certification authority, the DID of the governed party being certified, and the DID URL of the precise type of certification being awarded.

For additional recommendations about integration of the ToIP Governance Stack with the ToIP Technology Stack, please see the [ToIP Governance Architecture Specification V1.0](#) and the [ToIP Governance Metamodel Specification V1.0](#).

## 13. References

---

*NOTE: References in this first public review draft (PR1) are provided inline as hyperlinks. In the second public review draft (PR20, they will be added as separate lists of Normative and Informative References.*

## About the ToIP Foundation

---

Founded in May 2020 with 27 original founding member organizations, the ToIP Foundation has now grown to over 400 participating organizations plus over 100 additional individual participants. Our mission is to define an overall architecture for Internet-scale digital trust that combines cryptographic assurance at the machine layers (technology) with human accountability at the business, legal, and social layers (governance).

For more information about ToIP Foundation, please read our [Introduction to ToIP](#) white paper or visit our website at <https://trustoverip.org/>.

## Appendix A: Consolidated Requirements

For ease of reference, the following table consolidates all normative requirements in this specification. Each requirement is linked to the section in which it appears.

Req #	Description	Section
	<b>General ToIP Architecture Requirements</b>	
A.1	Implementers SHOULD ensure autonomy for ToIP Endpoint Systems.	<a href="#">7.1</a>
A.2	In a ToIP Endpoint System, the higher layers of the ToIP protocol stack MUST communicate with the lower layers via defined interfaces.	<a href="#">7.1</a>
A.3	A ToIP Intermediary System SHOULD be able to perform the functions of a ToIP Endpoint System for the purpose of routing enveloped messages using the ToIP Trust Spanning Protocol.	<a href="#">9</a>
A.4	The ToIP protocol stack in an Endpoint System MAY use the services of a Supporting System at any layer.	<a href="#">10.1</a>
	<b>ToIP Layer 1 Requirements</b>	
L1.1	If a ToIP Endpoint System includes Trust Support Functions, then those functions MUST be included at Layer 1 of the Endpoint System.	<a href="#">7.2</a>
	<b>ToIP Layer 2 Requirements</b>	

Req #	Description	Section
L2.1	A ToIP Endpoint System MUST communicate with another ToIP Endpoint System using the ToIP Trust Spanning Protocol.	7.3
L2.2	A ToIP identifier MUST be unique within the context in which it is used for identification.	8.2
L2.3	A ToIP identifier MUST be a verifiable identifier, i.e., verifiably bound to at least one set of cryptographic keys discoverable via an associated discovery protocol.	8.2
L2.4	A ToIP identifier SHOULD be a decentralized identifier, i.e., a verifiable identifier that does not require registration with a centralized authority.	8.2
L2.5	A ToIP identifier SHOULD be an autonomous identifier, i.e., a decentralized identifier that is self-certifying and fully portable.	8.2
L2.6	A ToIP identifier SHOULD support rotation of the associated cryptographic keys for the lifetime of the identifier.	8.2
L2.7	A ToIP identifier MAY also support rotation to an entirely different ToIP identifier that can be cryptographically verified to be a synonym of the original ToIP identifier.	8.2
L2.8	A ToIP identifier SHOULD support the ability to: a) associate the identifier with the network address of one or more ToIP Systems that can deliver to one or more Endpoint Systems under the locus of control of the ToIP identifier controller, and, b) if desired by the controller, enable that association to be discoverable.	8.2
L2.9	The ToIP Trust Spanning Protocol specification MUST define how to construct and format messages that are cryptographically verifiable to have the following four properties: (1) Authenticity: the message was sent from a sender who has control over the ToIP identifier. (2) Integrity: the contents of the message transmitted by the sender are received by the recipient without modification. (3) Confidentiality: the contents of the message are only accessible by authorized parties. (4) Privacy: the contents of the message are bound to conditions of usage agreed to by the parties	8.3

Req #	Description	Section
L2.10	In a ToIP Endpoint System, an implementation of the ToIP Trust Spanning Protocol MUST support authenticity and integrity.	8.3
L2.11	In a ToIP Endpoint System, an implementation of the ToIP Trust Spanning Protocol MAY support confidentiality and privacy.	8.3
L2.12	The ToIP Trust Spanning Protocol MUST enable the composition of higher-level Trust Task Protocols (such features as co-protocols).	8.3
L2.13	The ToIP Trust Spanning Protocol MUST support extensible message schema.	8.3
L2.14	The ToIP Trust Spanning Protocol MUST support resolution of ToIP identifiers to: a) the network addresses of receiving Endpoint Systems, and b) any required cryptographic keys.	8.4
L2.15	The ToIP Trust Spanning Protocol MUST support transport of messages via ToIP Layer 1 interfaces.	8.4
L2.16	The ToIP Trust Spanning Protocol MUST support delivery of messages to the Layer 2 interface of the Endpoint System of the ultimate receiver of the message.	8.4
L2.17	The ToIP Trust Spanning Protocol MUST support delivery of messages via Intermediary Systems.	8.4
L2.18	The ToIP Trust Spanning Protocol MUST support confidentiality with regard to the metadata required for message routing.	8.4
	<b>ToIP Layer 3 Requirements</b>	
L3.1	A Layer 3 Trust Task Protocol MUST communicate either over the Layer 2 ToIP Trust Spanning Protocol or over another Layer 3 Trust Task Protocol for all communications related to trust establishment between Endpoint Systems.	7.4
L3.2	A Layer 3 Trust Task MAY use other protocols, but only for other purposes (since short-circuiting Layer 2 when establishing trust with other Endpoint Systems would undermine the trust guarantees of the ToIP stack).	7.4

Req #	Description	Section
L3.3	A Layer 3 Trust Task Protocol intended to communicate private data SHOULD support Confidentiality and Privacy.	7.4
	<b>ToIP Layer 4 Requirements</b>	
L4.1	Layer 4 Trust Applications MAY use any number of Layer 3 Trust Task Protocols.	7.5
L4.2	If a Layer 4 Trust Application does not use a Layer 3 Trust Task Protocol, it MUST communicate with other Endpoint Systems using the Layer 3 Trust Spanning Protocol	7.5
L4.3	A Layer 4 Trust Application MUST support any ToIP-defined Trust Affordances relevant to that application.	7.5

## Appendix B: Consolidated Views of the ToIP Technology Stack

The ToIP Technology Architecture Task Force has spent many hours discussing how to produce consolidated views of ToIP architecture that are both relatively easy to understand but still technically accurate. In the end, we agreed that no single diagram is sufficient. Rather different views of the architecture should be taken together to see the whole picture. In this appendix we present several of these views — and [we invite feedback](#) on others that might be helpful.

### Functional Hourglass View

Figure B1 is a view of the types of functions that belong at each layer within a single Endpoint System as defined in this specification. It illustrates how the Hourglass Model is implemented as a single trust spanning protocol at Layer 2, with multiple trust support functions below and multiple supported trust task protocols above. It also shows one example (at the far right) of a specific category of Supporting Systems, in this case **verifiable data registries** (VDRs) upon which an Endpoint System can rely as external sources of truth.

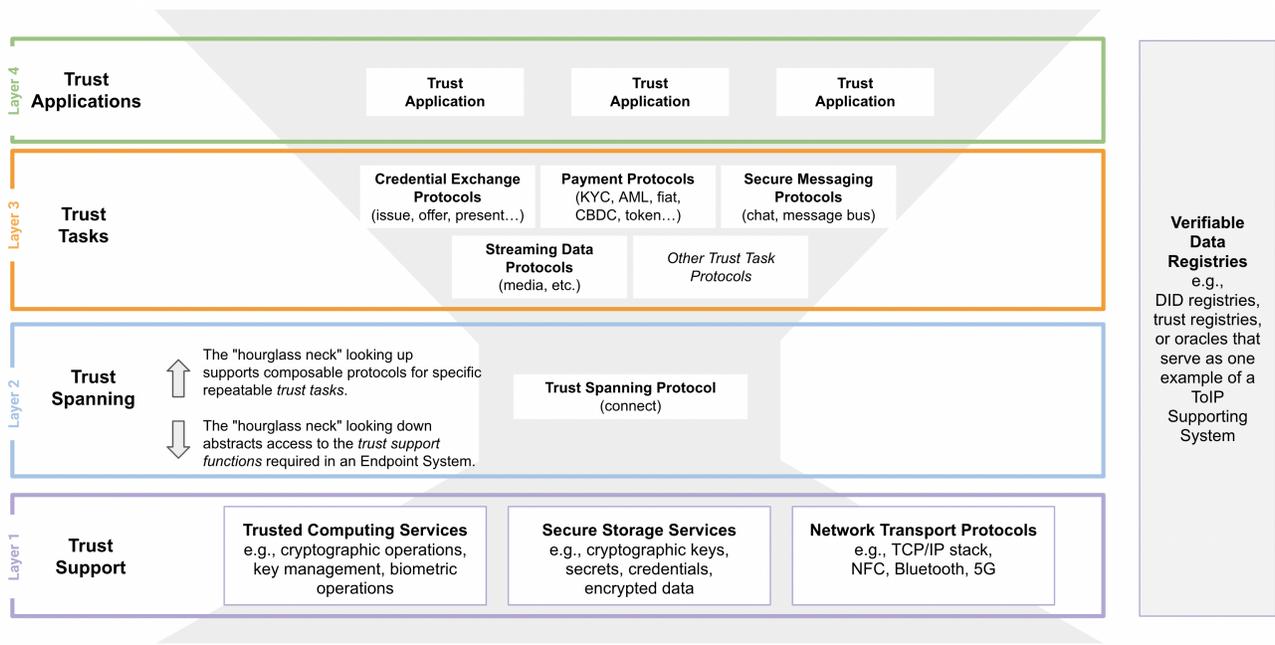


Figure B1: A layer-by-layer view of functions within an Endpoint System (also showing verifiable data registries as one type of adjacent Supporting System)

## Sphere-of-Influence View

Figure B2 builds on Figure B1 by identifying those technical capabilities that fall within the purview of ToIP’s technical architecture and those that are outside that boundary and thus do not need to be governed by ToIP component specifications.

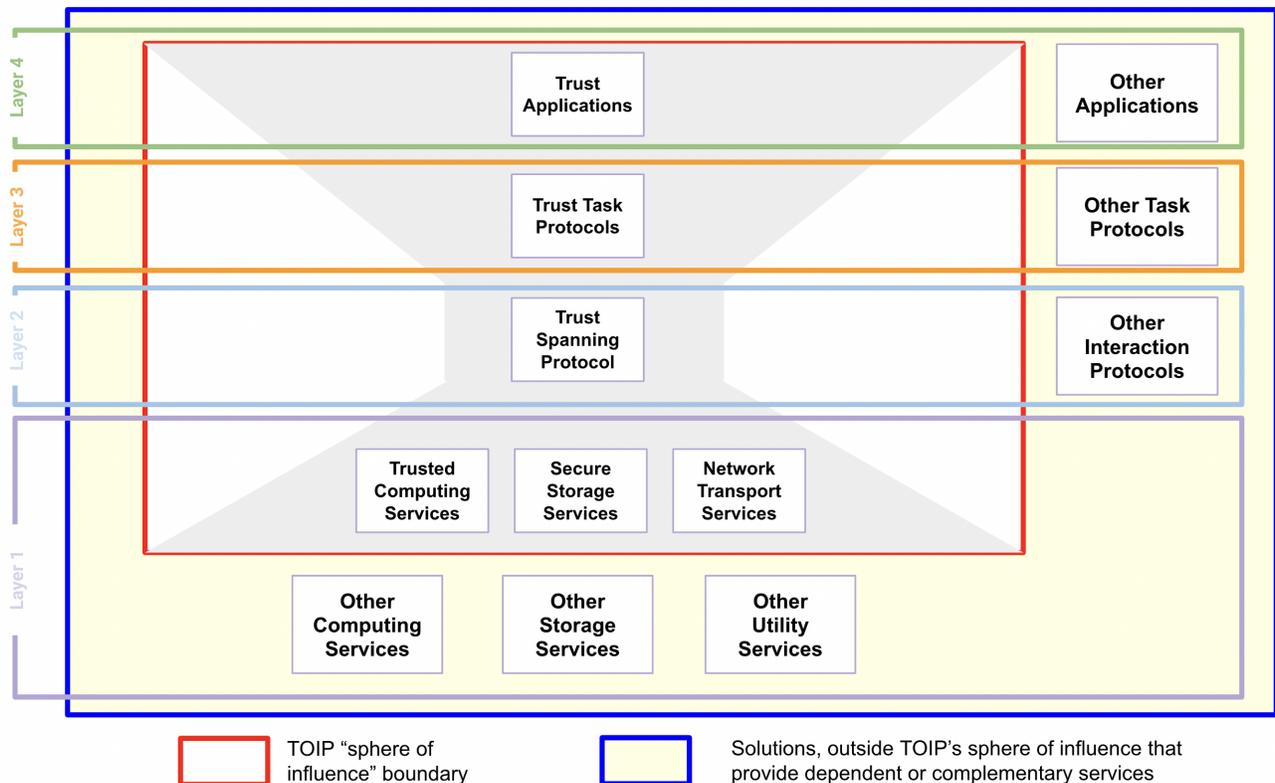


Figure B2: A view of the ToIP that shows what is inside and outside ToIP's "sphere of influence"

This view shows how the logical capabilities and components identified in the functional hourglass view can align with dependent solutions that are not governed by the requirements of the ToIP stack. For example, a DID resolver functioning at Layer 2 in an Endpoint System may call a DID ledger functioning as a verifiable data registry. While the DID resolver interface is a ToIP Layer 2 function, the DID ledger called by the associated DID method is a Supporting System that has its own resolution protocol as defined by the DID method.

## Interaction Pattern View

Figure B3 builds on B1 and B2 by showing the interaction patterns between two different Endpoint Systems as well as between an Endpoint System and a set of Supporting Systems (on the far right).

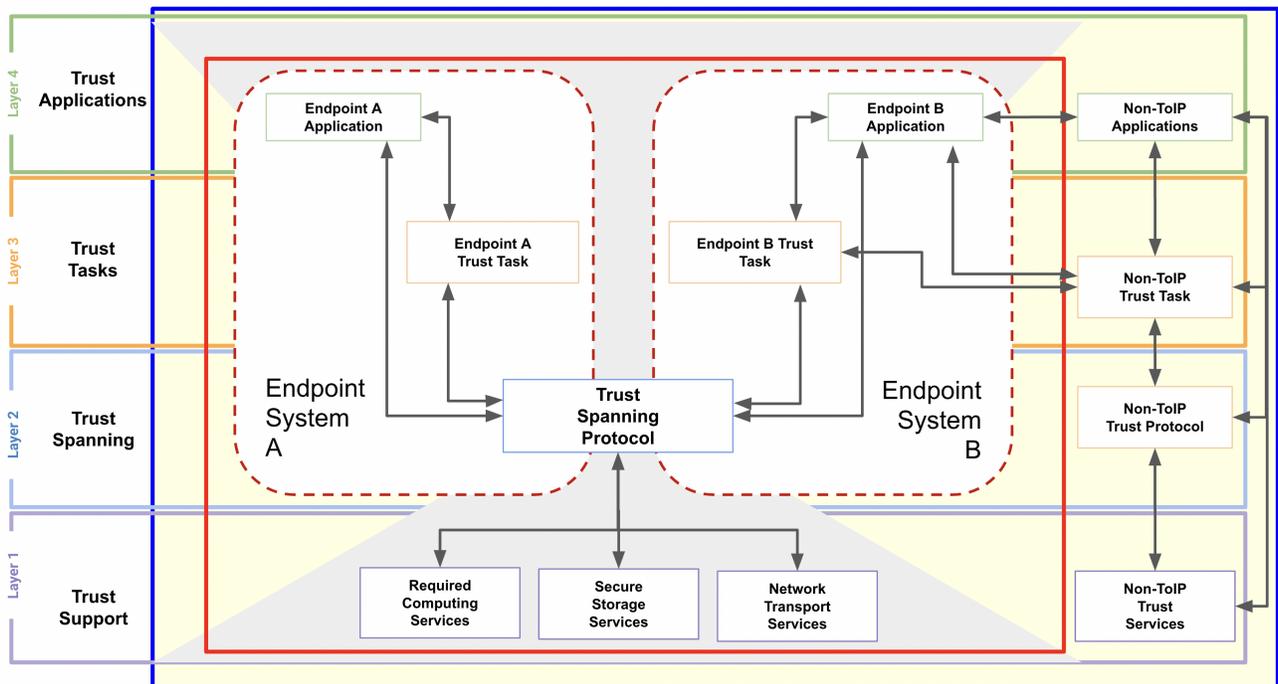


Figure B3: A view showing the interaction patterns both within and between two Endpoint Systems (as well as with Supporting Systems on the far right)

## Appendix C: Mapping of Existing Technologies into the ToIP Technology Stack

Just as the TCP/IP stack did not need to reinvent or replace existing local area networks, but instead added a new inter-networking layer to connect the existing networks, the ToIP stack does not reinvent or replace existing centralized or federated identity and PKI trust infrastructures. It adds a new inter-*trust* networking layer to connect the existing trust domains.

As this new layer of decentralized digital identity and trust infrastructure has been evolving, many individual pieces of the puzzle have been developed in parallel. For example:

- Over 120 different DID methods (see the [W3C DID Spec Registries](#)).
- At least a dozen different digital credential formats and signature schemes.
- At least three digital credential exchange protocols (DIDComm/Aries, OIDC4VC, W3C VC API)
- Several new decentralized payment protocols (e.g., [TBDex](#)).
- Multiple QR formats and out-of-band introduction (OOBI) protocols (e.g., [OASIS Secure QR Codes](#)).

From a ToIP perspective, all of these are potential component specifications.

- Some may fully meet the requirements of the ToIP Technology Architecture Specification without any modification.
- Others may require ToIP-conformant profiles to be written.
- Still others may require (hopefully small) revisions to meet ToIP requirements.
- A few (ideally very few) may need to be developed from scratch, either at the ToIP Foundation or another suitable industry organization.

Given the speed at which this new evolutionary branch of the Internet is evolving, the ToIP Foundation is maintaining a “mapping” of existing technologies and open standards into the ToIP stack in a separate living document called [Evolution of the ToIP Stack](#).

We highly recommend referring to this document to see the current mapping. The Foundation intends to publish an updated version of this document with each major development in the space (and no less frequently than twice a year).

1. With respect to this design goal, authenticity includes **message integrity**, i.e., a communication is not authentic if it has been tampered with in any way. 
2. Another important property of the architecture is **availability**. This is a concern with the design and implementation of operational deployments of the ToIP stack and should be addressed in the associated operational governance frameworks. 

