A Trust Layer for the Internet is Emerging

Toward a More Interoperable and Trusted Internet for Canadians

18DEC2023

Authors:

Darrell O'Donnell, Continuum Loop Inc.
Jacques Latour, CIRA

Table of Contents

List of Figures

Disclaimer and Author's Remarks

This report is published by Continuum Loop Inc. (Continuum Loop) and the Canadian Internet Registration Authority (CIRA). It represents the opinions of the authors and may not reflect the perspective of our respective employers.

We would like to thank our many reviewers for their comments and suggestions. In particular, thanks go to: Nicky Hickman, Cosanna Preston-Idedia, Drummond Reed, Christine Martin, Tim Bouma, Mathieu Glaude, Sam Curren, Marcus Ubani and Hadrien Seymour-Provencher.

Feel free to reach out for further discussions.

We look forward to hearing from you.

Cheers,

Darrell

Darrell O'Donnell
President & CEO
Continuum Loop Inc.

X/Twitter/etc: @darrello

Jacques

Jacques Latour
CTSO
Canadian Internet Registry Authority

Executive Summary

We are rapidly losing trust in the internet.

While the internet may provide secure communications and a plethora of walled gardens where we can operate, it doesn't help us build trustworthy relationships and systems.

Business leaders recognize that trust is critical to their success. Companies that are highly trusted outperform their peers (5%[1]), while those that experience trust-related incidents see major market capitalization change (20-56% decrease[2]).

Survey after survey shows a consistent and worrying trend - we are losing trust in the internet. Fakes, frauds, and forgeries dominate. Canadians, in particular, are losing trust faster than most of the planet. From 2019 to 2022, the percentage of Canadians who trust the Internet dropped from 71% to 57% (globally, the shift was from 74% to 63%)[3].

That trend needs to be reversed. Citizens and companies cannot build trust on top of an eroding system.

This report is about making the internet trustworthy again.

We identify that a **trust layer** is emerging that allows for an appropriate blend of **technical trust** (e.g. encryption and signing) and **human trust** (e.g. governance).

We know a blend of technology and governance is required to build this trust layer. Pairing technical and human trust allows us to create a trustworthy internet.

This emerging trust layer makes it possible, and preferable, to mutually authenticate each other. This means that parties both get a deep understanding of exactly who they are dealing with when they need that assurance. Knowing you are in a trusted digital relationship with an entity (e.g. your bank) while that entity knows it really is you on the other end creates a better way to build trust.

Specifically, this report is about two seemingly disconnected aspects:

- The emergence of a trust layer for the internet **-** where we can build trustworthy digital ecosystems on top of the internet. Blending technology with the governance we require for deep trust is critical to rebuilding trust on the internet.

[1] Deloitte. "The future of trust - A new measure for enterprise performance." *deloitte.ca*, https://www2.deloitte.com/content/dam/Deloitte/ca/Documents/deloitte-analytics/ca-en-future-of-trust-pov-cdn-aoda.pdf. Accessed 13 November 2023.
[2] ibid., p. 7.
[3] Ipsos. "Trust in the Internet." *ipsos.com*, https://www.ipsos.com/sites/default/files/ct/news/documents/2022-11/Trust%20in%20the%20Internet%2C%20Nov%202022.pdf. Accessed 6 October 2023.

- How **trust registries** help anchor trust points to bring technology and human trust into play. We discuss trust registries and the systems of record that provide authoritative answers for the digital ecosystems they serve. An additional **registry of registries (RoR)** concept is explored as it creates the web of connections we need to create an interoperable trust layer for the internet.

There is work to be done here. Businesses can begin to lay down foundations for which they can build trust on. More exploration is needed, but we know the broad shape of the terrain. Work is required, and much of the path is understood.

1. Introduction

This report started as a refresh of Continuum Loop's 2019 "The Wallet Report[4]." We were considering looking at what has changed to accompany the annual webinars we've been giving.

We realized that the wallets and agents mentioned in the 2019 report, while being good tools, are just building blocks. The more important thing is how we can build truly trustworthy systems.

Our focus shifted quickly to dig into what we now call "the emerging trust layer" and how we can re-establish trust on the internet.

The recent release of the World Economic Forum's "Reimagining Digital ID,"[5] where we had the privilege to contribute, stated that "the internet lacks an ID layer"[6]. While we agree with this notion, we would make a different statement:

> **…the internet lacks a trust layer.**

While we can build systems that have reasonably good security, we can't just pick up some tools that allow us to create truly trustworthy systems.

We can see small examples of highly trusted systems on the internet today, but they are an exception. While they may offer technical interoperability, they create closed environments. We know these as walled gardens.

It is time to create an interoperable, deeply trustworthy internet.

We know that there are a few things required - and this report will dive into the details:

- **Technical Trust** - we must have systems that don't cut corners that create limits to how trustworthy the systems built on them can be. We must know that our digital interactions are sufficiently confidential, authentic, and private.
- **Human Trust** - we need to bring our human constructs - the social and legal mandates in our lives, the formal governance, and the operations that keep us believing in each other. To date, this aspect has been utterly missing from the internet in any real way.
- **Registries** - we need to anchor what we can so the technical side of the ecosystems we build can automatically do the right thing. Knowing who we are working with is critical - except when we don't want/need to know them.
    - **Trust Registries** anchor ecosystems, giving solid answers to the human trust questions that get us towards being truly trustworthy.
    - A **Registry of Registries (RoR)** helps create a web of connected ecosystems, wiring up the existing disparate trust registries.

---

[4] O'Donnell, Darrell. "The Current and Future State of Digital Wallets." *continuumloop.com*, 2019, https://thewalletwars.s3.amazonaws.com/The-Current-and-Future-State-of-Digital-Wallets-v1.0-FINAL.pdf. Accessed 24 October 2023.
[5] World Economic Forum. "Reimagining Digital ID." *weforum.org*, June 2023, https://www3.weforum.org/docs/WEF_Reimagining_Digital_ID_2023.pdf. Accessed 25 October 2023.
[6] Ibid., p. 6.

2. A Trust Layer is Emerging

As we see more and more official information and services migrate online, trust failures are becoming increasingly apparent. High-assurance information is hard to trust when you can't determine who is authoritative and who is a fraudster. What is real, and what is fake? Introducing generative AI complicates whether someone is actually who they claim to be.

Fundamentally, the lack of a consistent way to establish trust means we lack certainty in our digital interactions.

We cannot reliably confirm the parties' authenticity nor ensure our conversations' confidentiality and privacy. We also can't easily or consistently prove the authenticity of the parties we work with.

To some extent, we know that the internet can be trusted, but only so far…

> **"We didn't focus on how you could wreck this system intentionally [when designing the internet]."**
> - Vinton Cerf [9]

As the internet evolved, key trust-related decisions were made, delayed, or ignored. These decisions - and the lack thereof, mean that the internet is missing key capabilities that help establish, build, and maintain trustworthiness across a changing landscape.

Post Snowden[10], the internet has been working hard to encrypt everything, ensuring no over-privileged eyes can see clear text information or metadata. But the internet community (e.g. IETF) has been focused on encrypting things, not signing them. The difference here means that though content may be encrypted, we

Domain Names and Certificate Authorities Combine to Create Limited Trust.

One of the only systems we have to establish authenticity combines the Domain Name System (DNS) and Certificate Authorities (CAs). It worked well to a point.

Think of the green lock symbol in your browser. A few things have happened when you see it:

- The domain name you see has a certificate.
- That certificate contains the domain name, which is mapped to the IP address of the server you are on.
- That certificate also contains cryptographic keys that are used to encrypt the communication channel using transport layer security (TLS).

This allows some basic trust to be associated with the communication channel. Still, it isn't nearly as trustable as desired - as there are deep flaws:

- The other end (you, if you are browsing the web) is utterly unknown to the website. This is why you have to log in with a username and password so often. This flaw renders two-way trust difficult and inconsistent.
- The domain name may not be the domain you are looking for. This wasn't as true in the beginning days of the internet[7], but now it is.
- There is no definitive way to know that the domain you are hitting actually represents the entity you mean to interact with[8].

[7] In the early days of the internet, getting an SSL certificate was very difficult. It required sharing a lot of information (e.g. corporate registrations) and relatively high fees to get a basic SSL certificate. At that point, say prior to 2015, having an SSL certificate was a good proxy to ensure that a site was who they purported to be. With certificates becoming low-cost or free this proxy lost most of its value.
[8] Edelman, Ben, and Tyler Moore. "Typosquatting." *wikipedia.org*, 2023, https://en.wikipedia.org/wiki/Typosquatting. Accessed 27 September 2023.
[9] Cerf, Vinton. *The Fate of Online Trust in the Next Decade*. 2017. *pewresearch.org*, https://www.pewresearch.org/internet/2017/08/10/the-fate-of-online-trust-in-the-next-decade
[10] Wikipedia. "Snowden effect." *wikipedia.org*, 2023, https://en.wikipedia.org/wiki/Snowden_effect. Accessed 25 October 2023.

don’t have the required authenticity - for the parties involved, nor for the information being exchanged.

Understanding how information is exchanged is critical to understanding what the emerging trust layer can become.

2.1. Digital Interactions

The internet, at its core, is all about digital interactions. It enables two or more parties to communicate and exchange information. It’s really that simple.

These digital interactions have enabled so much. They seem simple, but once we consider a few questions, things immediately get complicated and complex:

- What information is being exchanged?
- Who are the parties involved in an interaction, especially if there are more than two involved?
- What are the roles of each party?
- Do they need the same information?
- Can each party see all the information exchanged or just a part of the information?
- What have the parties agreed to that can/cannot be done with the information exchanged? How is that agreement discussed and enforced?

In the world of the internet, exchanging data is a common practice. These are the digital interactions that happen trillions of times a day. We can think of a digital interaction as having three fundamental building blocks:

- **Relationship** - information exchange occurs within the context of a relationship between two or more parties. This relationship serves as a foundation for the exchange and sets the stage for how data will be handled.
- **Information** - data that has been processed and contextualized, turning raw details into meaningful insights.
- **Channel** - information is exchanged over a communication channel that connects the parties.

This leads to a simple definition:

- **Digital Interaction** - *information* exchanged over a *channel* in the context of a *relationship*.

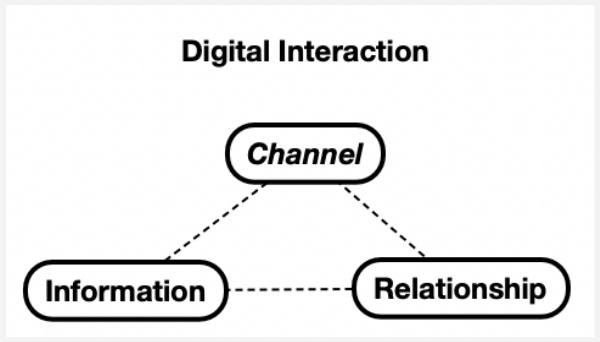

***Fig. 1** - A Digital Interaction - Information exchanged in the context of a relationship.*

The question of trust arises when we attempt to determine if a specific digital interaction should be trusted enough to satisfy the parties involved, in other words, whether they can consider their interaction trustworthy.

That is where things have gone badly wrong on the internet. We don't have built-in mechanisms for establishing trust.

The problem lies in our limited understanding of the authenticity of information and the authenticity of the parties engaged in the information exchange.

We are missing the trust. We don't have **trusted interactions**. We can't fully trust the:

- **Information** - information can be tampered with or observed in many contexts.
- **Relationship** - we can't be sure that the parties are authentic.
- **Channel** - communication is over a channel that never had security and authenticity built in.

2.2. Low-Trust to High-Trust

On today's internet, we need to do a fair bit of homework to get beyond the low-trust that exists (and is sliding to no trust). Trust in the internet is dropping rapidly. From 2019 to 2022, only a three-year period, an IPSOS[11] study showed:

- Trust in the internet dropped almost 15% globally and almost 20% in Canada from 2019 to 2022.
- Concern about privacy is increasing globally, with 79% of people concerned about privacy.

To begin the shift back to a high-trust environment, we need to understand what is required to shift our digital interactions from the low-trust level they are at to become trusted interactions - where trust-enabling capabilities are baked in.

That starts with understanding an interesting triad that requires some tradeoffs.

[11] Ipsos. "Trust in the Internet." *ipsos.com*, 2022, https://www.ipsos.com/sites/default/files/ct/news/documents/2022-11/Trust%20in%20the%20Internet%2C%20Nov%202022.pdf. Accessed 6 October 2023.

2.3. Confidential, Authentic, and Private

In order to trust a digital interaction, we need assurance that our interaction is confidential, authentic, and private. These three dimensions are critical to understand:

1. **Confidentiality**: Are the contents of an information exchange protected so only authorized parties have access?
2. **Authenticity**: Are the parties to an information exchange able to verify that the information originated from the correct sender and has not been tampered with?
3. **Privacy**: Will the expectations of each party concerning the usage of shared information be honoured by the other parties?

We call these dimensions the **CAP** (confidentiality, authenticity, and privacy) dimensions for trust. However, by following best practices, we can maximize these dimensions. We can also recognize when we make trade-offs (like sacrificing privacy for authenticity - a common tradeoff) and assess if it aligns with our trust requirements.

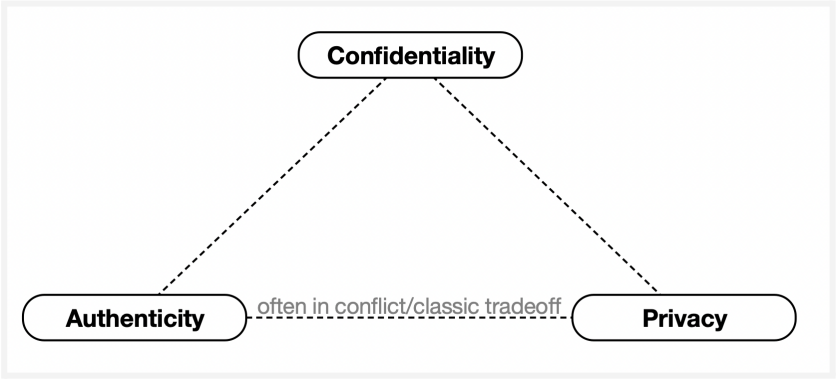

***Fig. 2** - The CAP dimensions for trust.*

2.4. Technical Trust and Human Trust

Establishing technical trust for a trust layer is necessary but not sufficient. While the technical trust level can be high, without governance being interwoven, the best one can hope for is to improve confidentiality. That means we are leaving out improvements we need for authenticity and privacy.

We need to understand how to answer critical questions, such as:

- How do you know that a procurement officer has the legal authority to sign a particular contract?
- How do you know that a particular device is owned and operated by a particular corporation? What happens when that device is sold or lost and it still has a connection to your system?
- How do we combat deepfakes and deal with AI-generated content? Can video camera manufacturers and processors provide a solid chain of custody that lets us know, with reasonable certainty, that what was shown is what actually happened? Initiatives like

the Coalition for Content Provenance and Authenticity (C2PA) have been created[12] to solve this - but they require a trust layer for full viability.

For that, we need to consider a different kind of trust—human trust. To clarify the difference:

- **Technical Trust** - can I prove that the data came from a particular source and that it hasn't been tampered with?; and
- **Human Trust** - can I trust the other parties that I am working with (i.e. are they really the source that I think they are)? Have they followed the required rules and processes that we need? Do they have legitimacy to a level that I can rely on?

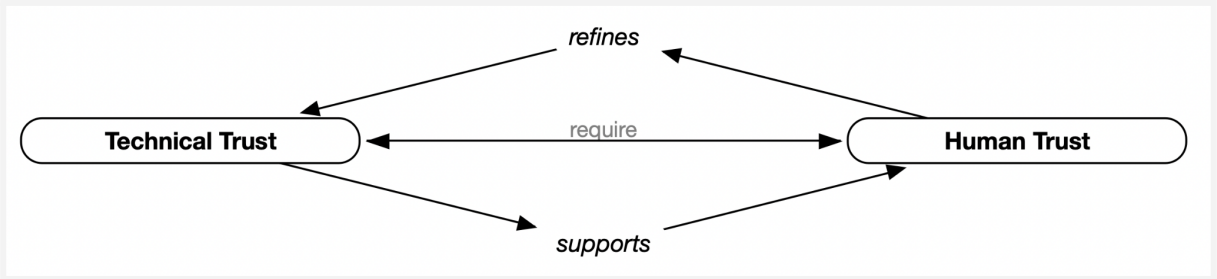

***Fig. 3** - Interplay between Technical and Human Trust*

These concepts build on each other:

- Human trust is supported by technical trust. At the technical trust level, flaws in design or implementation can restrict how much human trust can be enabled. If the technical trust of an interaction is low, there will be fundamental limits to how trusted an interaction can be.
- Technical trust is refined by human trust - Given the CAP requirements determined by the interaction context, the technical trust will be configured to maximize the particular dimensions required. This is often where the authenticity and privacy tradeoffs are codified.

The human trust side comes from governance, which is a broad area. We'll cover that shortly, but first, we will look at the technical side.

2.5. Payload, Protocol, Parties, and Purpose

On a high level, in our trusted interaction model, we need to understand **what** is being exchanged - the trusted information - as well as **how** the parties are making the exchange. Additionally, we need to know **who** is doing the exchange. This "what," "how," and "who" simplifies how we evaluate what needs to be interoperable. Even more important may be **why** the interaction is happening.

Earlier, we discussed that our digital interactions are about exchanging information between parties over a channel. This maps out to more technical terms:

[12] Coalition for Content Provenance and Authenticity. "C2PA Founding Press Release." *C2PA.org*, 22 February 2021, https://c2pa.org/post/c2pa_initial_pr/. Accessed 6 November 2023.

- **Payload** - covers the information shared in an interaction. The payloads must be understood at both a structure (schematic) level and a meaning/context (semantic) level. Different ecosystems use similar-sounding ideas but have utterly different meanings. Assumptions over structure or meaning can have catastrophic results.
- **Protocol** - covers the back-and-forth related to information exchange (the payloads) required to conduct a trusted interaction. This will include the used protocols and exclude consequences of successful and failed interactions.
- **Parties** - the entities that are exchanging information.

To trust an interaction, we need to know the following:

- That the **payload** hasn't been tampered with and came from a trusted source; we can trust the information.
- That this trusted information is being exchanged using a **protocol** that respects our CAP requirements; and
- That the **parties** in the relationship are who they say they are.

In addition to the what, how, and who (of payload, protocol, and parties), we can also assign a reason - or **purpose** to the interaction. **Purpose** provides the reason, the "why", we are conducting an interaction.

This means we can have a trusted interaction. When we are exchanging trusted information over a trusted channel, in the context of a trusted relationship, we have a trusted interaction.

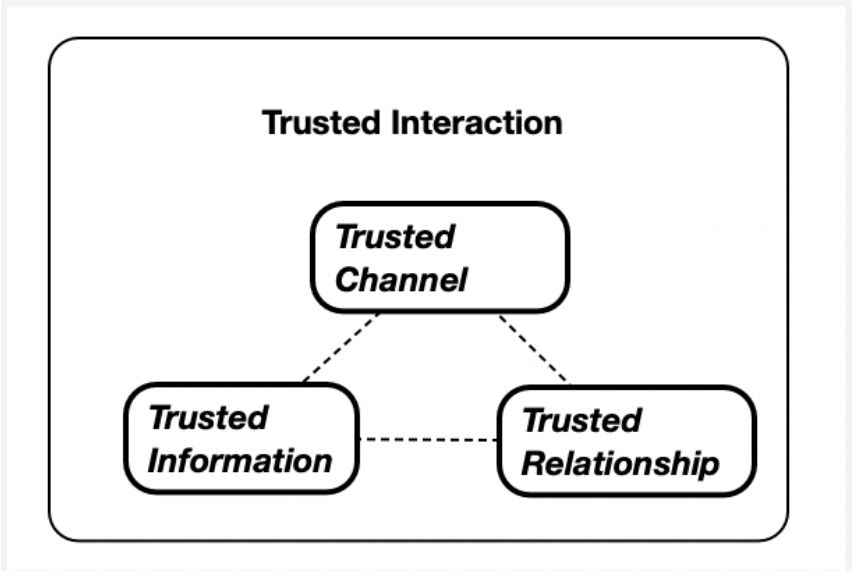

***Fig. 4** - Components of a Trusted Interaction*

This trusted interaction concept means we can begin to understand how a trust layer can be used over the internet. However, this requires that we blend both the technical trust and human trust that we discussed earlier.

To do that, we need to try and simplify some incredibly complex things. We need to realize that we operate across so many domains at a time that looking at the requirements for human trust gets blurry. This blurring is because technology is insufficient to meet our technical trust and human trust requirements.

We need governance for that, and governance is messy. Further, we need our technology to at least not undercut our governance needs and, ideally, to embody them.

Separating and joining the governance and technical stack is critical.

2.6. Governance and Technical Stacks

A significant reason the internet is becoming less trusted and less trustworthy is that the internet was created largely to solve technical problems. While governance comes into play in some areas (e.g. DNS), the internet is principally a technical construct.

Efforts to bring governance to bear have been difficult, but progress is being made. In particular, the Trust Over IP (ToIP) Foundation[13], hosted by the Linux Foundation[14], is making headway.

Considering that technology and governance must work together is a core tenet of the ToIP Foundation, which is building out a model that creates a dual governance and technology stack, the ToIP Model[15].

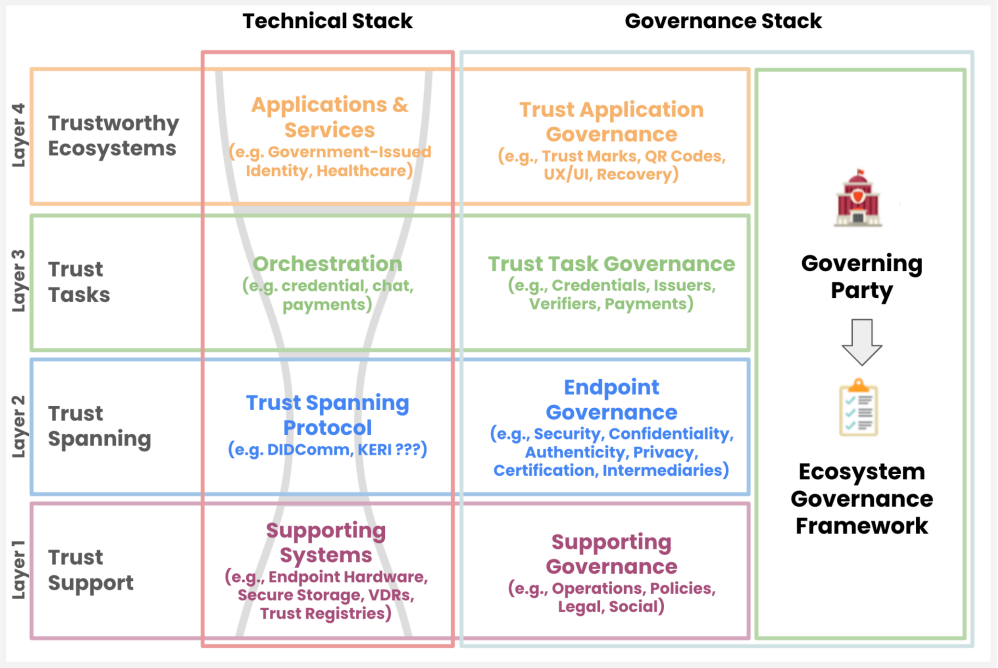

Fig. 5 *- ToIP Foundation 4-layer dual stack*

[13] www.trustoverip.org
[14] www.linuxfoundations.org
[15] Trust Over IP Foundation. “The ToIP Model.” *trustoverip.org*, 2022, https://trustoverip.org/toip-model/. Accessed 4 October 2023.

The general concept is that technology and governance influence each other and need to consider how they reinforce. Further, the technology and governance are layered. These layers are where we need to focus for the remainder of the discussion about the emerging trust layer. In particular, two of the ToIP layers are most relevant:

- **Layer 4 -** the **Trustworthy Ecosystems** layer creates the conditions for simplifying how we view the rich and complex governance requirements for real-world use of a trust layer. More on that in Section 3.
- **Layer 2 -** The **Trust Spanning** layer creates the technical and governance conditions to meet the CAP requirements. It allows for the confidentiality, authenticity, and privacy required for any ecosystem.

While Layer 1 (Trust Support) and Layer 3 (Trust Tasks) are also critical, they are less important for explaining what the emerging trust layer requires. This is because both of these layers are highly dependent on which ecosystem they are supporting. The specifics change with each ecosystem, but the patterns are the same:

- **Layer 1 - Trust Support** functions are the fundamental building blocks and resources used by an ecosystem and higher layers.
- **Layer 3 - Trust Tasks** are the actions conducted in an ecosystem. They create a reusable and composable framework to support generic ecosystem needs. Consider these as ingredients. Each ecosystem provides recipe(s) and uses the ingredients.

For more detailed information beyond what is covered in this report, we recommend exploring the ToIP Technical Architecture Specification[16] as it describes how an overall trust architecture can be built on top of the internet in a four-layered architecture.

In the next section, we will discuss some aspects of the envisioned **trust spanning layer**, which is analogous to the TCP/IP and UDP layers that played a pivotal role in creating the modern internet. We will focus on how the trust spanning layer can lay the foundations for high levels of trust, and we will explore how the CAP dimensions play a crucial role in facilitating this enhanced trust.

2.7. Trust Spanning Layer

Unlike the current internet, the approach in designing the trust spanning layer intentionally allows us to establish technical trust without compromising human trust. The design is focused on ensuring that the CAP requirements for a trust layer can be met.

We aren't going to go deep into the technical specifications for a trust layer, but we can point to solid work done in the design of the Trust Spanning Protocol (TSP) as described in the 2023 progress report[17]. That resource goes deep into explaining things. Our focus here is about how we use the TSP to create a trustworthy internet.

---

[16] Trust Over IP. "TechArch/spec.md." *GitHub*, https://github.com/trustoverip/TechArch/blob/main/spec.md. Accessed 27 September 2023.
[17] Trust Over IP Foundation. "Mid-Year Progress Report on the ToIP Trust Spanning Protocol." *trustoverip.org*, 31 August 2023, https://trustoverip.org/blog/2023/08/31/mid-year-progress-report-on-the-toip-trust-spanning-protocol/. Accessed 27 September 2023.

In essence, the TSP focuses on establishing technical trust and allowing for human trust to be interwoven. The TSP provides a robust technical framework for CAP, particularly for information exchange. Within this framework, it manages identifiers, cryptography, and messaging to facilitate the maximization of CAP.

The Trust Spanning efforts maximize Confidentiality and improve the ability to have high levels of Authenticity and Privacy. This effort allows us to establish a very high level of technical trust.

Establishing a foundational layer for technical trust requires approaches for handling data and communication integrity, ensuring data provenance, applying technical controls and performing verification and validation (and more). At a fundamental level, this requires an understanding of how some key aspects of our digital interactions are set.

None of the CAP dimensions can be achieved solely through technical means. They require human trust to understand the inevitable tradeoffs between the CAP dimensions. The ToIP Model dual stack (technical and governance) helps here. Technical trust is embedded in the ToIP Technical Stack, while human trust is embedded in the Governance Stack.

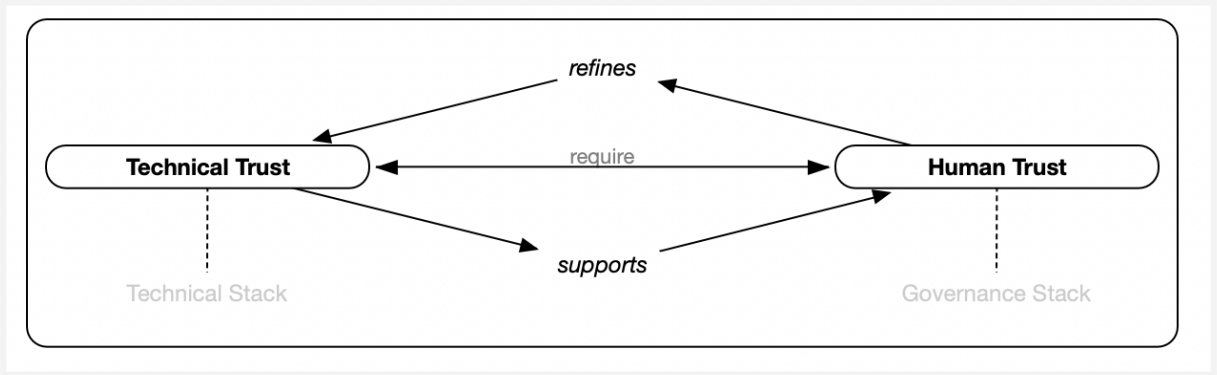

Fig. 6 *- Technical Trust, Human Trust, and ToIP Stack*

Further, when examining the tradeoffs of the CAP dimensions, it is important to understand that they exist on a continuum. Meeting the required CAP dimensions involves varying degrees of both technical and human trust.

Confidentiality, as an example, is largely a technical trust issue with some minor human trust adjustments. Conversely, privacy is primarily a human trust issue supported by technical trust. Authenticity sits in between, depending on the context of an interaction. For a detailed discussion about how CAP impacts both technical trust and human trust levels, see Appendix C.

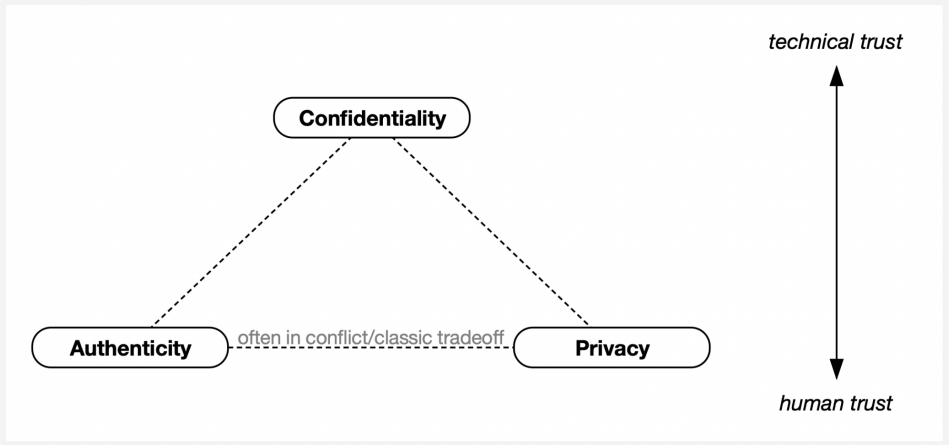

***Fig. 7** - CAP Dimensions on the Technical-Human Trust Continuum*

It is critical to recognize that the TSP efforts alone cannot fully enable the conditions for high trust. While allowing for human trust, the TSP does not directly establish it.

Human trust and, particularly, its governance aspects must be considered.

2.8. Human Trust, Governance, and Ecosystems

We have been using the term “human trust” relatively frequently. We mostly mean governance in the context of ecosystems, so let’s discuss that. We won’t be going too deeply into detail as there are plenty of resources to look at there. We’ll cover some key concepts, though, as they are critical for our discussion:

- There are potentially multiple ecosystems at play, and the parties to a trusted interaction may have very different roles in each ecosystem.
- Governance can provide us with a comprehensive understanding of who the other party is and what rules we agree to operate under and abide by.

We have used the term “governance” relatively loosely to this point. It is time to define how we use it more tightly. We are looking specifically at how parties are interacting in a digital ecosystem. In this case, effective ecosystem governance involves creating an environment that fosters innovation, ensures fair competition, protects against harmful practices, and promotes sustainability.

While the intricacies of governance may be out of scope for this report, we need to understand the following:

- We already have many governance constructs in place - for government identity documents, industry-specific information sharing and credentialing of professionals. This means that in high-assurance cases, we can generally find an answer to “Who has the list of authoritative sources?”. Although the process of compiling, formalizing, and keeping these lists up-to-date can be resource-intensive, the required information is accessible and must be.

- We operate in many different ecosystems. When a bank needs to see government identity documents, they aren't running the government identity ecosystem - they are participating in and relying on it.
- When checking sources, we need to understand what authority they state that they operate under. That points us to the governance constructs they operate under and the ecosystems they operate within.

"I forgot my card" - Multiple Ecosystems At Play

As an example, going into a credit union without your debit card to conduct a transaction may involve two different but related ecosystems. For simplicity, we'll call these:

- Financial ecosystem - follows the rules and regulations, both external (nation-state defined) and internal (to the bank), that define modern financial systems.
- Government Identity ecosystem - the state-issued identity credentials (e.g. driver's license, passport, birth certificate, etc.) that can be used to identify a person.

The credit union is deeply involved in the Financial ecosystem - they are a financial institution. They are less involved (though they may think they can control personal identity) in the Government Identity ecosystems - they rely upon that ecosystem as opposed to being deeply embedded.

When a credit union member walks in without their debit card and needs to do something (e.g. take out some cash), the credit union:

- Looks at the rules for their *Financial* ecosystem first. They see that they need to fully identify the member before allowing cash to be withdrawn.
- That's where they rely heavily on the *Government Identity* ecosystem to accurately identify the person they can link into their system.
- Once they have satisfied that the person in front of them matches the government identification provided, they will likely go through many more internal checks, as required by the *Financial* ecosystem they are part of.

The key point here is that we have two distinct ecosystems, and the Credit Union and Person (credit union member) see them differently.

While this report isn't about governance, knowing that many (most) interactions occur in the context of one or more ecosystems is quite important. Depending on the trust level required for a trusted interaction, the ecosystems involved may be complex. What is critical is understanding - and believing - that the related governance requirements are being met.

Each ecosystem will have its own governance - both formal and informal. We are using the terms:

- **Social mandate** - an informal set of governance constructs that come from informal sources that are usually socially imposed.
- **Legal mandate** - a formal set of governance that comes from formal sources such as regulation and legislation that apply to a particular ecosystem. Additionally, any formal ecosystem-adopted legal constructs (e.g. a contractually enforced ecosystem governance framework) can be applied.
- **Business mandate** - the various operational processes, procedures, best practices, etc., that are expected in a particular ecosystem.

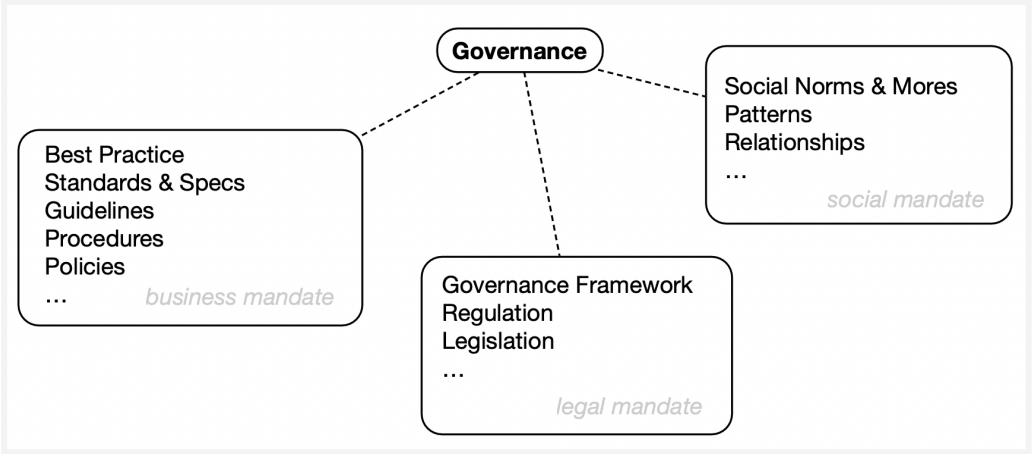

***Fig. 8** - Governance and Mandates (Institutional and Social)*

2.8.1. Trusted Relationships - Are You Authoritative?

As mentioned in the prior section, one fundamental governance and technical question is:

“How do we know the other party is who they say they are?”

For trusted interactions, this question is critical. Technology can’t solve it alone - this requires governance - human trust.

We need to know that the issuer of a government (or some other) credential is legitimate and authorized (or authoritative) to do what they are doing. Otherwise, we can’t trust them or their data.

So, when you operate in a decentralized manner, where do you get an answer that says the issuer of a government identity credential is genuinely who they say they are and that they are authoritative?

The simplest approach to knowing if an entity is who they say they are is looking to a trust registry.

But first, we need to understand that there is another critical thing that can be provided by a trust registry.

2.8.2. Defining Human Trust In Digital Terms

The hard part of using the terms “human trust” and “technical trust” is that while we want to codify everything, the human part is difficult to quantify and qualify. All to often, this “squishiness[18]” of the human trust is left out of a system.

The good news is that some simple constructs can be created to provide technical (digital) answers to the hard-to-define elements of governance and human trust.

Getting answers to some of these key questions can be done with a trust registry as well.

[18] Darrell, one of you author’s, uses this term in business for the hard to quantify and qualify elements. It is used to acknowledge areas where we know ambiguities remain.

3. Trust Registries

Trust registries have become a clear requirement for any digital ecosystem, particularly those that are decentralized. Counter-intuitively, they provide the information needed to really anchor the trust layer.

A trust registry provides answers that are considered authoritative for an ecosystem.

More specifically, a trust registry is a system (e.g. a system of record) that answers questions within the context of its authoritative ecosystem, including:

- Questions about entities and their authorizations within the ecosystem.
- Configuration data specific to that ecosystem.

Answering these two questions allows systems to join an ecosystem and get access to the information that is needed.

Examples can include:

- A list of valid government identity document issuers.
- A list of licensed professionals.
- A list of financial institutions connected to a global payment messaging system (e.g. SWIFT).
- Configuration data that maps out what levels of assurance, authorization types, etc., are used in the ecosystem.

You'll note that the word "list" is prevalent.

A trust registry looks like a list externally. Internally, it may have rich functionality that enforces governance and operations (e.g. adding a new financial institution to a national register is a very rich process). However, externally, the information is quite simple (e.g. a new financial institution is just another element in the list).

Getting the inputs required to make trust decisions becomes challenging without a trust registry. This difficulty often leads to poor decisions due to the unavailability of authoritative information.

Some definitions of the term "trust registry" may be set for various reasons and be more restrictive than the above definitions. For example, the Pan-Canadian Trust Framework[19] provides the following very identity-centric definition:

> *The purpose of a trust registry is to provide participants of a Digital Identity Ecosystem the means to verify that Ecosystem digital participants are trustworthy[20]. Just as a Verifier needs to know they are dealing with the right person, as outlined in the Verified Person*

[19] DIACC. "Pan-Canadian Trust Framework." *diacc.ca*, https://diacc.ca/trust-framework/. Accessed 25 October 2023.
[20] The word "trustworthy" isn't used in the same way in this report, as the decision about whether a trust registry is "trustworthy" depends on who is consulting the trust registry (i.e. they decide what they deem trustworthy).

component and that their Credential is valid, the Verifier and the Holder need to know that the organization that is issuing the credential is valid. If an Issuer is listed in a trust registry, this indicates to interested parties (e.g., Verifiers and Holders) that an Issuer can be trusted as an authoritative provider of credentials. Digital Identity Ecosystems and their associated trust registries use a Trust Framework (such as the PCTF) to define how Issuers, Verifiers, Holders, and Digital Wallets should or must operate to be considered trustworthy.[21]

Other ecosystems may have wildly different needs; examples can include:

- Global industry initiatives (e.g. Towards Sustainable Mining[22] initiative) where key players need to be understood on a global basis but are managed regionally/nationally.
- The decentralized reputation of entities and their authorizations managed in a decentralized autonomous organization (DAO). Traversing a blockchain and constantly building up data is good for provenance and underlying belief, but it doesn't scale well.

Our role in various ecosystems varies wildly. Depending on our activity, we may be deeply involved or only incidentally touching upon an ecosystem. We need to be able to anchor our interactions. This is where trust registries really help.

3.1. Trust Registries Anchor The Trust we require in the CAP model.

- How can we connect? This question concerns the technical detail necessary to perform a trusted connection.

Answering those questions is the primary purpose of a trust registry.

However, both questions are challenging to answer in a one-size-fits-all manner. The simplest, yet technically least satisfying, response is: "It depends."

We ask these questions because we are attempting to determine if they should be considered trustworthy.

For that, we require more context.

Determining whether an entity is authorized to perform a specific action is highly dependent on context, as is understanding how systems can connect (to achieve what purpose?) is context-specific.

This is where the concept of ecosystems becomes crucial. We often operate within different ecosystems, each with its own set of tasks and objectives.

[21] DIACC. "PCTF Trust Registries Component Overview PCTF13." *diacc.ca*, 1 March 2023, https://diacc.ca/wp-content/uploads/2023/03/PCTF-Trust-Registries-Component-Overview_Draft-Recomendation-V1.0r.pdf. Accessed 25 October 2023.
[22] TSM. "About." *Towards Sustainable Mining*, https://tsminitiative.com/about. Accessed 25 October 2023.

3.2. Ecosystems and Trust Registries

When you ask the questions above in an ecosystem context, answers become much easier to deliver:

- Is this entity authorized to do something **in an ecosystem**?
- How can we participate **in an ecosystem**?

Within a well-governed ecosystem, a trust registry can answer these questions.

But let's back up and discuss what we mean by "ecosystems" and how they help frame things.

3.2.1. We Operate in Many Ecosystems

Ecosystems are widespread and often overlap. Identifying existing ecosystems that can assist with tasks your systems shouldn't handle alone is important. Many ecosystems share responsibilities, with some managing specific aspects and relying on other ecosystems for different functions.

Let's consider an example: a situation where two parties transfer funds between two countries.

The following simplified[23] list of ecosystems applies where parties between two countries are moving funds:

- Government-Issued Identity Ecosystem - each nation has its governance and procedures for issuing identity credentials (physical and/or digital).
- Financial Institution Ecosystem - each nation has its own ecosystem for financial ecosystems.
- International Funds Transfer Ecoystem - a supranational-level ecosystem coordinating fund transfer.

These ecosystems are linked when funds move from Party1 (in Nation1) to Party2 (in Nation2). Many individual things can be illustrated here:

- Governments issue identity credentials to people (and organizations often) that are used in their daily lives. Each country implements things differently, but there is a general pattern that is the same.
- The Financial Institutions depend on the Government-Issued Identity Ecosystem, but they don't control it.
- The Financial Institutions must comply with their nation's national (and sub-national) governance as defined (and implied) in their nation. One standard piece that is relevant here is the compliance rules (e.g. KYC, AML/CFT) that apply in each country. Part of the national governance likely enforces specific standards that also apply internationally.
- The International Funds Transfer Ecosystem has its own governance framework, which combines international standards with the rules of the participating nations. They have authority at the international level but can only exert influence at the national level.

[23] There are many other systems and ecosystems at play. This example is being used for explanatory purposes. Additional detail would take away from the explanation.

- The International Funds Transfer Ecosystem will have requirements for identifying the parties that the Financial Institutions would handle (using the Government-Issued Identity Ecosystem in their respective countries). It is key to understand that there is no direct link to the Government Issued Identity Ecosystem, but the information will flow - via the Financial Institution Ecosystems.

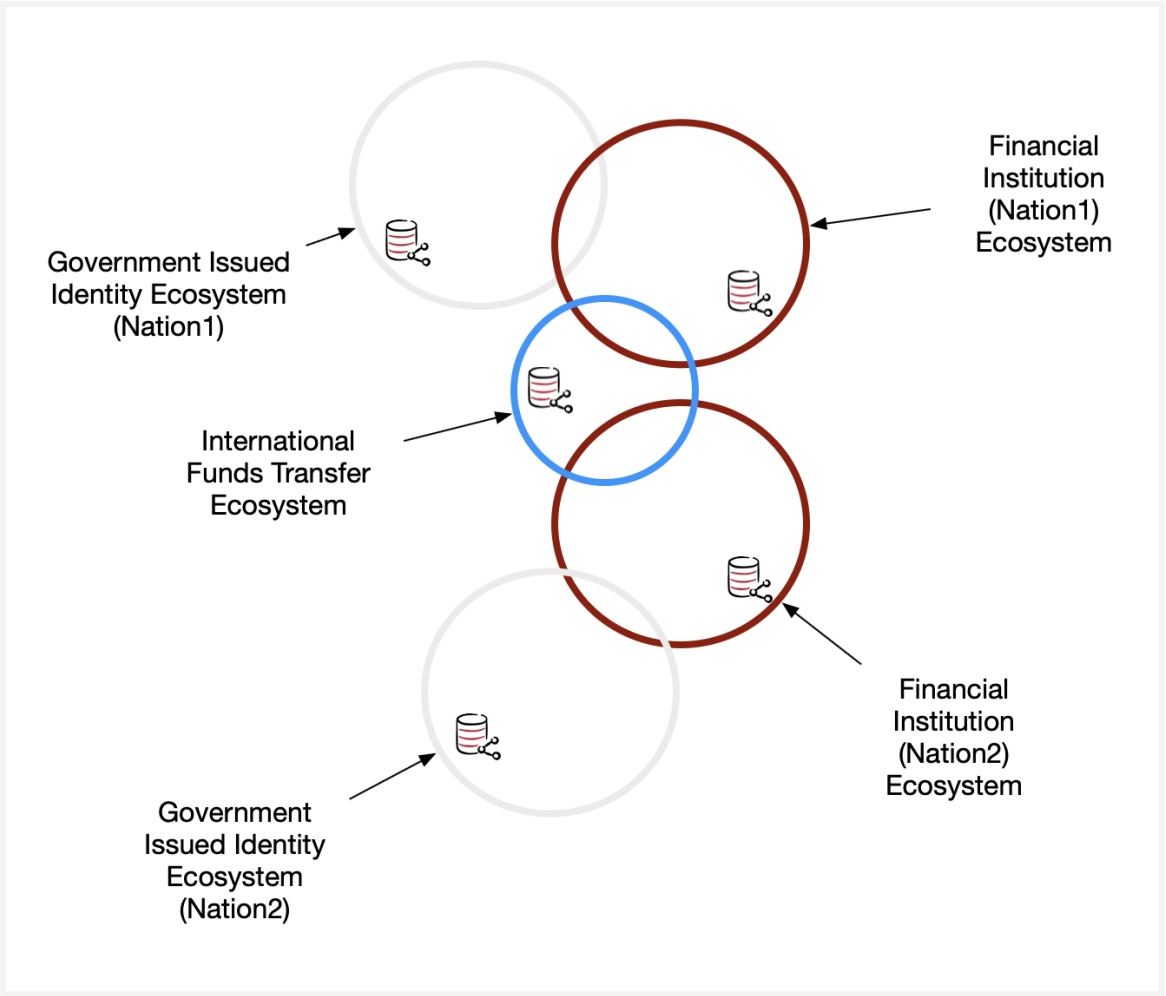

Fig. 9 *- Interconnected Ecosystems in International Fund Transfers*

Each of these ecosystems would likely have one or more trust registries. That's important to understand. A bank in Nation1, as part of its duties defined in the Financial Institution Ecosystem (for Nation1), would likely rely upon one or more trust registries to get a list of authorized government identity credential issuers. Similarly, when a bank needs to know if they can connect with a bank in another country, they would consult a trust registry in the International Funds Transfer Ecosystem as it maintains the system of record in its ecosystem.

3.2.2. Trust Registries Anchor Ecosystems

Each of the ecosystems we operate within has its unique characteristics and attributes. They may overlap highly, but when we dig into them - they have a discrete set of players and governance. Trust registries are one tool that helps anchor things together. An ecosystem's governance roles and activities can be represented in and served by a trust registry.

These trust registries help us ground the trusted interactions that, as we saw earlier, lead to a trust layer for the internet.

3.3. Trust Registries Expose Governance

When navigating the intersection of human and technical trust, we inevitably have a collision between the human interpretation of governance and the software that makes rigid decisions. The best we can hope for at this point is to obtain actionable answers.

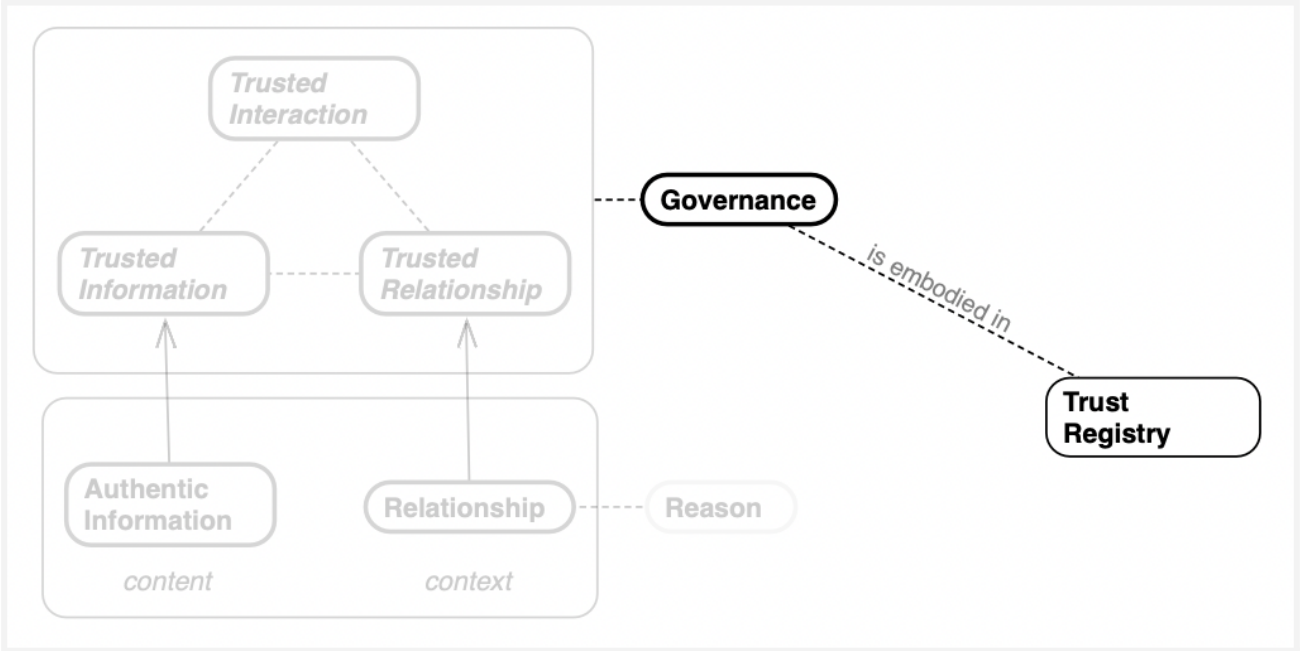

Fig. 10 *- Trust Registries and Decision-making*

The current efforts to define trust registries, and more specifically, how outside systems can get the answer to governance questions, have landed on two sets of simple technical questions that a trust registry can answer:

- What do I need to understand your system? This is the configuration information that systems need to connect and operate in an ecosystem.
- How can I find out about a particular entity's authorization in an ecosystem?

These questions are asked in the context of an ecosystem that the trust registry anchors.

The value of a trust registry is clear, but if they have to be created from scratch, we have a lot of work to do before we can use them.

3.4. We Don't Need to Build Them All - Trust Registries Are Everywhere

Many folks in the digital trust space are concerned about building the raw infrastructure to support the Authenticity problem. How can we build trust registries and populate them with data about entities holding various forms of authority? The task sounds daunting.

The good news is that most basic trust information is broadly available. Many systems that maintain the information we need to establish trust have existed for several decades or longer (e.g., corporate registries, professional certification bodies, governments, etc.).

The challenge lies not in creating new sources of information but in accessing the existing systems that hold the answers we need. In other words, identifying an authoritative source isn't the main difficulty; the work is in establishing a connection to a system that can provide the necessary answers. The good news is that making the connection isn’t that difficult.

The reason for this vast difference in accessibility and availability is simple: the existing systems haven’t needed to answer questions digitally until now. Further, they didn’t know exactly what questions would be asked.

Fortunately, steps are being taken to make that less of a problem.

A specification for the protocol to communicate with trust registries is under development.

3.5. Trust Registry Protocol

The concepts shared about trust registries to this point have been somewhat abstract and instructional. The process must be concrete to allow trusted interactions to operate on an internet scale.

We need a protocol so systems can interrogate a trust registry simply and consistently.

Work on a standardized protocol for communicating with a trust registry is underway at the ToIP Foundation. The Trust Registry Task Force is working to create a new version of the Trust Registry Protocol (TRP).

Past efforts (TRP version 1) created a protocol overly focused on only Verifiable Credential use cases. It answered three specific questions:

- Does the trust registry recognize an Issuer as Authoritative for a Credential Type under a particular Ecosystem Governance Framework?
- Does the trust registry recognize a Verifier as Authorized for a Presentation Type under a particular Ecosystem Governance Framework?
- Does the trust registry acknowledge another trust registry under a particular Ecosystem Governance Framework?

The draft TRP (TRP version 2) creates more generic capabilities, allowing richer governance constructs to be represented. Further, it explicitly considers how the systems of record that provide the trust registry capabilities can support the TRP.

The TRP version is focused on the two types of questions mentioned in Section 3.3:

- Configuration information that provides the base information needed to connect into an ecosystem. Examples include:
    - What **Authorizations** are managed by the systems that comprise the trust registry?
    - What **Levels of Assurance** are supported by this trust registry?
    - What other data (metadata) can be shared about this trust registry?
- Dynamic questions about authorizations of entities in an ecosystem. Examples:
    - Does a particular **Entity** have a particular **Authorization** under the ecosystem governance framework that governs the trust registry?
    - Is another trust registry **recognized** by this trust registry?

The protocol allows systems to adopt the TRP natively or create simple bridging software to access legacy systems.

The identifiers used by each entity in a trust registry are unique. While the initial work on trust registries began at ToIP with decentralized identifiers (DIDs), centralized identifiers are also supported. This allows PKI (e.g. X.509) systems to be integrated.

The TRP is intended to be used in at least two key modes (see Figure 11):

- **Bridged** - where systems already have digital sources, it is simple to create a bridging service that uses the system of records database and/or APIs via the TRP.
- **Native** - systems of record may directly build in support for the TRP, as it standardizes and simplifies the queries that external systems require.

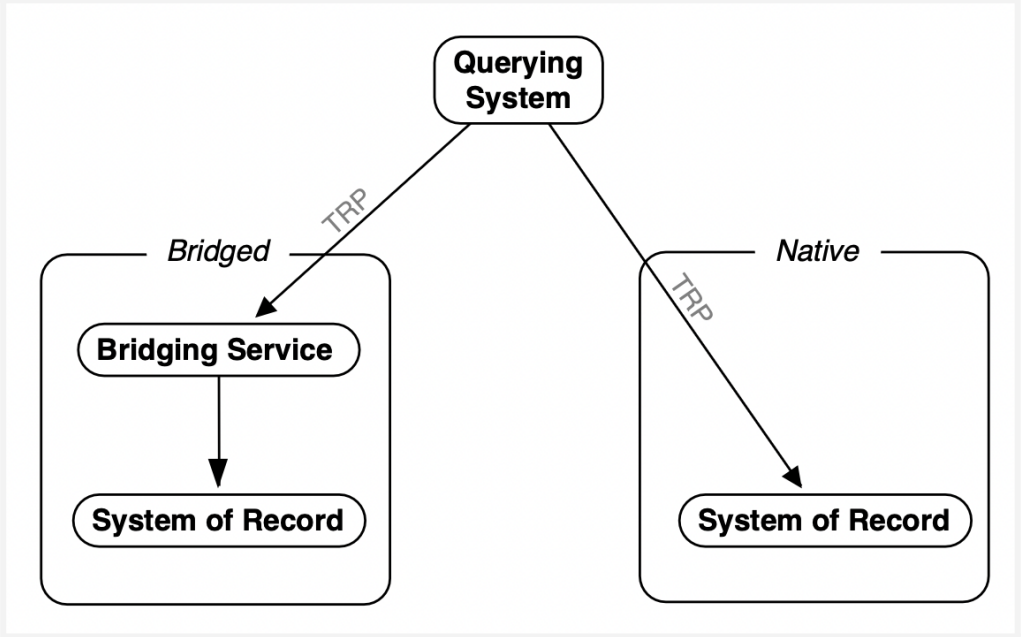

***Fig. 11** - Integration Patterns: Native or Bridged TRP Approaches*

3.6. Trust Registries are Not (Necessarily) Centralized

The web3/decentralized community seems to think that any concept of centralization is anathema. This is quite ridiculous, as points of centralization will always occur. Further, there are logical points

of centralization. The question is where centralized authority is optimal. This depends a lot on what problem you are solving[24]. Consider the following:

- When we need to understand national identity for international travel, we naturally end up at the nation-state (countries issue travel passports) and perhaps international (ICAO coordinates the passport scheme; more on that later).
- When I want to understand your reputation, it's a complex mix of centralized (e.g. Twitter/X, Facebook, identity documents) and decentralized (informal - and likely more valuable - networks, people vouching for you, and more) elements.

It is all contextual, and a trust registry can in most contexts.

A trust registry and the TRP have no opinion on where the point of centralization is. A trust registry can support answers backed entirely by centralized, distributed, and decentralized systems. The hard work is about understanding your context and applying the right technical and human design.

Initially, that context is likely about answering basic trust questions that lend themselves to more centralized solutions. However, even those may be decentralized, or as a dear friend taught your author - they may follow the concept of Subsidiarity[25]. Put simply, subsidiarity means locating decisions where the best information and context for decision-making can be. The same industry may look wildly different in different jurisdictions.

Interesting work is going on at IETF[26] for further reading about centralization and decentralization.

3.7. Trust Flywheel

As trusted interactions prove more and more valuable, the role of trust registries becomes more apparent. Trusted interactions drive value up, and trust registries drive costs down.

Once you see the value of a trust registry, finding more trust registries will inevitably increase both the quantity and quality of your trusted interactions. They create a flywheel: more trusted interactions increase the value of trust registries; more trust registries increase the quality and quantity of trusted interactions, driving the volume.

[24] Nottingham, M. "Centralization, Decentralization, and Internet Standards." *ietf.org*, 9 July 2022, https://www.ietf.org/archive/id/draft-nottingham-avoiding-internet-centralization-05.html. Accessed 28 September 2023.
[25] Wikipedia. "Subsidiarity." *wikipedia.org*, 2023, https://en.wikipedia.org/wiki/Subsidiarity. Accessed 25 October 2023.
[26] Nottingham, M. "Centralization, Decentralization, and Internet Standards." *ietf.org,* 30 August 2022, https://www.ietf.org/archive/id/draft-nottingham-avoiding-internet-centralization-14.html. Accessed 28 September 2023.

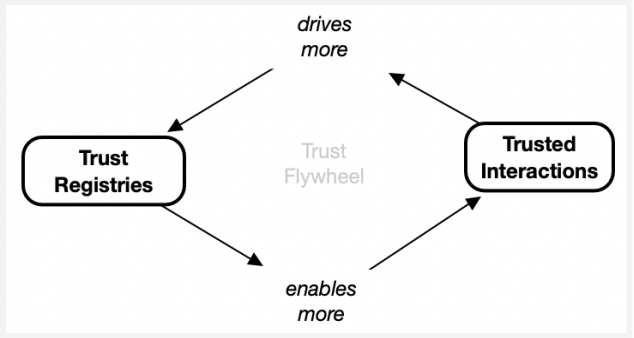

***Fig. 12** - Trust Flywheel*

This flywheel pattern drives the need to find more trust registries. Those that already exist and those trust registries that need to be connected.

But how do we find these trust registries to work with?

3.8. Discovery

There are at least three different patterns at play when we are looking for trust registries:

- **Self Declaration** - Issuers should be clear about which trust registries they are registered in to assist in discovery.
- **Informal Approach** - system builders will gradually learn about relevant trust registries to help them work.
- **Formal Governance** - where trust registries are built into an ecosystem, the EGFs will reflect this.

While these approaches are feasible, they don't scale well. They are all basically a centralized approach to a decentralized and distributed problem.

The number of potential trust registries can easily hit thousands and even millions. With the vast number of potential trust registries, it is infeasible for any organization to maintain solid enough connections to all the required information sources. The list cannot reasonably be centralized.

However, every participant needs to know where they can go to get answers.

What is needed is a place where known registries can be discovered.

In a fully hierarchical world, we could look for a root registry.

However, the internet and how the world operates are distributed and decentralized. There is no common locus of control for the broad range of trust registries that are needed.

We need something decentralized yet globally useful.

A "registry of registries" concept has emerged as a solution to this scaling and discovery problem.

CIRA, our partner in this report, has funded research into the Apex Registry[27] concept. Using the term “apex” has sparked a strong reaction, which has helped drive discussion and learning about what is meant and needed. The concept of a registry of registries has evolved from this work and interactions in other communities.

[27] IDLab. “Pan-Canadian Trust Registry Community of Practice.” *Digital Identity Labratory*, 7 August 2023, https://www.idlab.org/en/pan-canadian-trust-registry-community-of-practice/. Accessed 25 October 2023.

4. Registry of Registries

When there is a need to inventory and discover if there are trust registries that can be considered, a specialized trust registry is required. One way to meet this need is a registry of trust registries - a registry of registries (RoR).

The idea behind a RoR is for some authority (e.g. a national registry or ecosystem registry) to assert that they are aware of - and curate - a list of known registries that fit their mandate.

4.1. Internet Scale:

As we look at it globally, we begin to realize that the potential number of trust registries is beyond any organization's ability to scale. Consider the following breakdown by country, NGOs, and trade associations:

- Countries - there are 193 member states of the UN. Each country is responsible for numerous industries and sectors.
- Non-Governmental Organizations - there are millions (over 1.5 million in the US[28] alone) of NGOs that may play a role in managing trust registries.
- Trade Associations - thousands of trade associations operate at global, regional, national, and sub-national levels. Trade associations set standards and define ecosystems in many industries. They maintain the data needed for trust registries.

The system here is not hierarchical - at least not in a consistent way. The UN does not "control" nation-states. We have reached a point where some companies effectively transcend national boundaries - a *technopolar* moment, as described by Ian Bremmer[29].

This means we need to handle a distributed and decentralized set of trust registries. There are many ways to examine how various ecosystems can distribute their governance and tools that will drive different patterns. For this report, we will focus on two key approaches that are likely to help:

- **Nation State RoR** - similar to country-code top-level domains (ccTLD) in the DNS realm, we see patterns where both domestic and international needs are served well in some key ecosystems.
- **Ecosystem Specific RoR** - Some industries are best handled on an ecosystem/industry basis as they have governance that may be global, regional, national, or subnational.

Most value-generating ecosystems will depend on other ecosystems that may be run on a different basis. As an example, consider the Towards Sustainable Mining Initiative's hierarchical nation

[28] Bureau of Democracy, Human Rights and Labor. "Non-Governmental Organizations (NGOs) in the United States." *state.gov*, 20 January 2021, https://www.state.gov/non-governmental-organizations-ngos-in-the-united-states/. Accessed 11 October 2023.
[29] Bremmer, Ian. "What is a technopolar world?" *gzeromedia.com*, 30 August 2023, https://www.gzeromedia.com/ai/what-is-a-technopolar-world. Accessed 11 October 2023.

basis (e.g. auditors for carbon credits likely need to be accredited in the country they are working in). See the pullout (A Global Sustainable Mining Ecosystem) below for more details.

A Global Sustainable Mining Ecosystem

Towards Sustainable Mining (TSM) "is a globally recognized sustainability program that supports mining companies in managing key environmental and social responsibilities."[30]. The TSM Initiative is a new program that has begun to set and enforce standards to improve the sustainability of mining globally. TSM is a voluntary, global program with 13 member countries and many multi-national companies as partners. The countries each administer their own regulations, legislation, and compliance, but they must meet the global TSM standards. The companies that are members often operate internationally. The TSM Initiative maintains global governance, while each member country would maintain its own trust registry (or registries).

The TSM ecosystem depends on organizations in nation-states. It would require trust registries (and RoRs) at both an ecosystem and national level:

- **Ecosystem** RoRs:
  - **Governing Organizations** - The TSM Initiative is global but, under its governance approach, delegates some responsibilities to national-level organizations. In Canada, the Mining Association of Canada is the responsible organization.
  - **NGOs** - the TSM initiative reports to non-government organizations as third parties. These NGOs are, by definition, outside of a nation-state structure.

- **National** RoRs:
  - **Auditors** - there are likely both ecosystem-based and nation-based RoRs. Many countries (e.g. Canada and the US) have accountant/auditor designations managed by regulated organizations at the province/state level. Other countries may manage these designations directly at various levels of government.
  - **Government Reporting** - the TSM initiative has mining companies reporting to various government bodies that have jurisdiction over their activities in the country.

4.2. National or Ecosystem-Based Registry of Registries

There is a common pattern that occurs on various geographical scales. In a particular ecosystem, the governance of an ecosystem may operate irrespective of geography. The common parts are more about the ecosystem (e.g. higher education) than the countries they operate in. In some industries, the approach doesn't fit a national registry approach - as they are more ecosystem-focused.

Several examples can be provided for registries that would make sense to organize on an ecosystem basis:

- **International Carbon Credit Accounting** - While nation-states will certainly have legislation and regulations that must be met, the coordination is transnational. An ecosystem body may maintain a global RoR, which points to various trust registries in many countries.
- **Global Health** - the World Health Organization already plays a coordinating role in various aspects. It may be the logical locus of control for some transnational health initiatives, likely pointing its RoR at the national trust registries that are run by the WHO member states.
- **International Payments** - the example of international fund transfer we discussed in Section 3.2.1 requires a transnational ecosystem approach to coordinate the cross-border transfers.

[30] TSM. "About." *Towards Sustainable Mining*, 2023, https://tsminitiative.com/about. Accessed 11 October 2023.

4.3. Registry of Registries Example - Two Ecosystems and Two Countries

Let’s consider two working examples in two countries. While similar, they expose slightly different patterns that are helpful to understand.

When Canadians and the international community have questions about authority, they don’t know where to start to determine if an organization is authoritative for something it is doing.

Today, even with paper and plastic credentials, we don’t really “know” that something is official. We can make good guesses, but…

- How do we know if a digital university degree credential is real?
    - Is it from an accredited university in Canada?
    - Is it from an accredited university outside of Canada?
- How can a hotel in Osaka, Japan, know that a Canadian government-issued photo ID is real versus fake? They need to know where to start - and starting with the issuer’s DID is relatively simple.

Let’s consider a couple of cases:

- Case 1 - Canadian applying for a professional designation in Canada
- Case 2 - Canadian students applying for a Ph.D. program in Japan

4.3.1. Case 1 - Domestic Education & Employment Credentials

A student who has graduated from a professionally accredited program has applied for their professional designation (e.g. P.Eng/ing or CPA) with the registrar for their profession.

The registrar must confirm that the degree presented meets their requirements and is from an accredited program. Further, they need to confirm that the degree was issued to the same person who is presenting a government identity document.

The manual process requires matching photo identity documents (e.g. driver's license, passport) to the person. Then, matching the person's name to the degree presented (name mismatches are a problem), and then confirming with the institution (i.e. calling them) that the degree/diploma is valid.

When a RoR and trust registries are in place, the whole process can be largely automated - and far less error and fraud-prone:

- The system confirms the digital credential issuer (DID) authenticity and also confirms the credential authenticity.
- The system then discovers the primary trust registry associated with the DID and ensures the authenticity of the issuer DID.
- The system then discovers any RoRs associated to the primary trust registry, in this case the system would be anchored in a Canadian RoR and can make sure that the issuer is indeed authentic and authorised to issue these credentials.
- Following this process, the system can confirm that the applicant holds a valid government ID
- And the system can also confirm that a valid educational institutions issued the degree by consulting the educational.
- Trust registries that they use and confirming accreditation with the appropriate provincial trust registry (educational institutions are provincially governed in Canada).

4.3.2. Case 2 - International Post-Graduate Education

A student has applied for a Ph.D. program in Japan. They hold undergraduate and master's degrees from two different Canadian universities.

The university registrar in Japan needs to know that the degrees are “real.”

Currently, it is difficult for the registrar. With valid universities being created and/or renamed officially and fake degrees[31] common enough for concern, where does the registrar start? The process here becomes complex and nearly byzantine quickly. Many universities have restricted the international universities that they will even consider - vastly limiting opportunities.

With a RoR approach, the registrar would:

- Same process as above,
    - Extract the issuer’s DID and affiliated trust registry.
    - Ensure the creds are authentic.
    - Ensure the DID is authentic with the trust registry.
    - Ensure Canadian RoR trusts the trust registry.
    - Ensure Japanese RoR acknowledges the Canadian RoR.
- Except the verifier in Japan would be anchored in the Japanese RoR, in this case, the verifier would need to ensure that the Canadian RoR trusted by the credential issuer is trusted by the Japanese RoR.
- Confirm government identity document - the registrar would likely hit a Canadian trust registry to confirm that an authorized entity issued the government-issued credential. The registrar would need to understand the country-level ecosystem but could defer trust to that ecosystem. Otherwise, they would need to maintain links to all of the valid government issuers in Canada. Depending on the country, there may also be a RoR where the registrar could start.
- Confirm validity of the degree - the registrar a country-level RoR to get a list of known universities (some countries would manage this nationally) or a list of the trust registries for each province and territory in Canada. This allows them to validate and verify the digital degree for the student instantly. The registrar likely will have other criteria to check (e.g. rankings of undergraduate schools), but they don’t need to worry that a degree was faked[32].

While these scenarios may seem relatively straightforward, both countries have complex systems for education and government-issued credentials. Neither country maintains a full awareness of how each country is organized. In Canada, for example, education and government issued identity are quite different:

- Education is largely managed at the provincial/territorial level - though there are exceptions where that is not true (e.g. Royal Military College falls under federal jurisdiction). Education is an area of churn - institutions are regularly created, renamed, or retired.

[31] Government of Canada. “False Academic Credentials.” *canada.ca*, 7 January 2019, https://www.canada.ca/en/public-service-commission/services/oversight-activities/investigations/summaries-investigation/summaries-investigation-fraud/18-19-03-false-academic-credentials.html. Accessed 25 October 2023.

[32] The case where an internal university employee fraudulently issues a degree is not solved here. However, it would provide evidence that could be used to adjust a university’s status in various trust registries.

- Government Issuers of Identity Credentials are largely provincial/territorial with exceptions:
    - Immigration, Refugees, and Citizenship Canada (IRCC) is a federal department that issues:
        - foundational identity documents for immigrants, refugees, and new citizens.
        - Additionally, they are responsible for issuing passports. The federal government regularly changes structure, and the responsibilities for immigration, refugees, and citizenship may be in different government departments.
    - The Department of National Defence issues driver's licenses.

The key here is that we have at least fourteen trust registries for each domain in our use cases (ten provinces, three territories, and one federal government). Once you add that universities, colleges, trade schools, and other institutions may be managed differently, you can quickly get to well over twenty sources of information. There are almost no international standards that cover global use.

Japan cannot be expected to know all these subtle aspects of Canada's education and government-issued identity ecosystems. Similarly, Canada can't be expected to manage information about the Japanese systems.

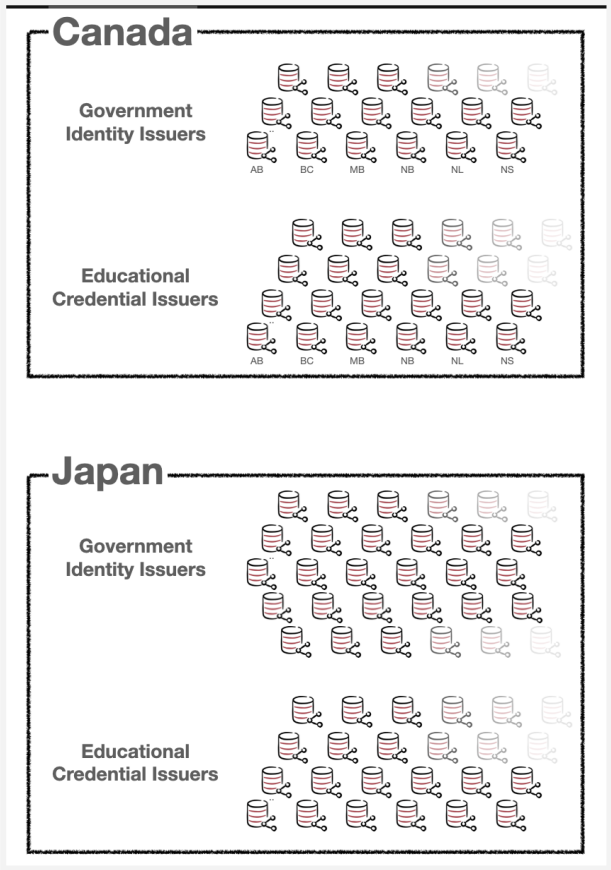

Fig. 13 *- Comparative Visualization of Trust Registries: Canada and Japan*

There is at least one notable exception to this lack of international standards - due to international work to recognize travel passports. In the example above, IRCC is the current department that issues passports in Canada - and the department name regularly changes. Japan doesn't need to

know that because the passport credentials are governed by the ICAO passport ecosystem[33]. This ecosystem has rich governance, technology, and a rigorous compliance process. While the Canadian issuing department may change, the ICAO requirements and usage don't. Interestingly, the ICAO passport system is basically a trust registry. Japan can use the ICAO passport credential schema and know that any issuing country's passport is in compliance.

The passport example doesn't necessarily help in our examples, though. The various university registrars may not know how to spot a fake or fraudulent passport. Border officials have the tools to do what others do not.

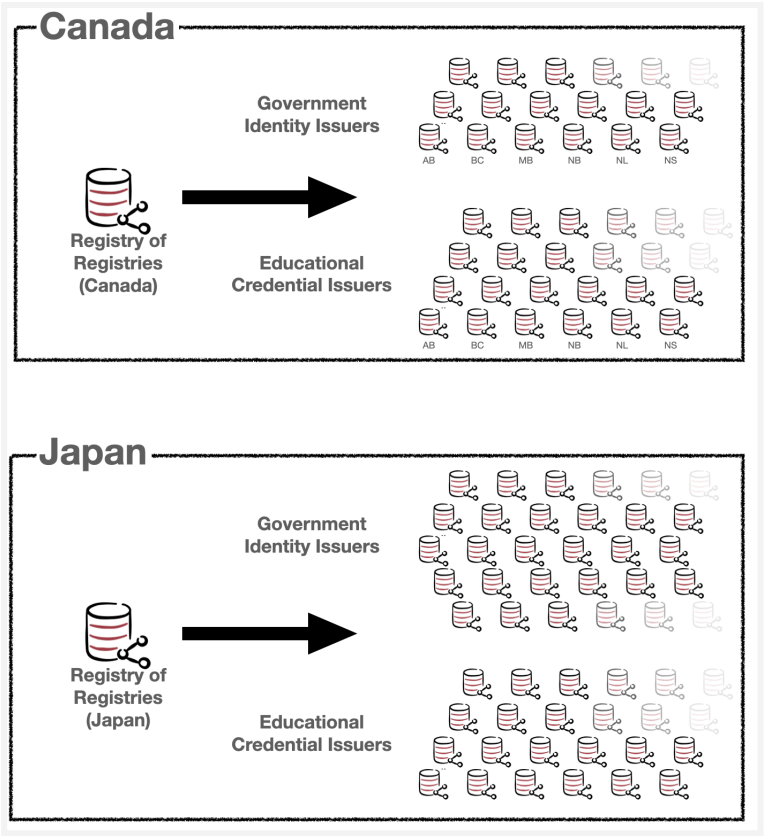

Fig. 14 *- Comparative Visualization of Registry of Registries: Canada and Japan*

4.4. DNS Parallel

The pattern of looking to a particular level - a country or ecosystem - is quite similar to DNS's top-level domain (TLD) pattern.

TLDs are managed and served in a decentralized and distributed way - much as these registry-of-registries will need.

[33] Government of Canada. "Features of Canada's new passport." *canada.ca*, 29 June 2023, https://www.canada.ca/en/immigration-refugees-citizenship/services/canadian-passports/new-passport-features.html. Accessed 25 October 2023.

4.4.1. Country-Code Top-Level Registry of Registries

The analogy best compares nation-state registries - a country code RoR. Another term could be a country code top-level registry (ccTLR).

On a national level, a ccTLR would be a RoR that serves as a gateway into a country's hierarchy of recognized ecosystems (e.g. education, identity credentials, finance). Any system that needs to understand the niceties of a particular country could start at the ccTLR. Queries interested in educational credentials could query the Canada ccTLR to get access to the trust registries that are managed provincially, territorially, and otherwise. This way, the querier would get consistent information.

4.4.2. Ecosystem Top-Level Registries

On an industry ecosystem basis, a similar approach could be taken. As an example, the Towards Sustainable Mining initiative could host its TLR in its own DNS records at tsminitiative.com.

The pattern similarity - from DNS to RoR - warrants some consideration. Many ecosystems are managed by a collaborative group. The TSM initiative mentioned earlier is managed by a community of interest under specific governance and uses a global domain (tsminitiative.com) as its main collaboration point. That domain could easily host RoR service for the TSM mining community.

4.5. Establishing a Registry of Registries Protocol

At the time of writing, there is an effort to review the Trust Registry Protocol (TRP) discussed in Section 3.5. The authors are involved in this effort and believe there is potential to merge the needs of trust registries and RoR concepts.

To support a registry of registry pattern, the existing TRP will need to be extended to allow one or both of the following:

- Querying for a list of acknowledged trust registries based on an ecosystem indicator (e.g well-known string or Ecosystem Governance Framework).
- Pass-through queries that allow querying multiple trust registries via a single RoR endpoint.

While other options may exist, the above have been floated as ideas in mid-2023.

Each ecosystem will require its own namespace for the trust registry types to be recognizable globally. Where namespace collision is a risk, or where the namespaces are globally relevant, these names should be coordinated with a body like ICANN[34] or IANA[35].

[34] https://www.icann.org/
[35] https://www.iana.org/

5. Conclusion

Canadians and the global community are witnessing a rapid decline in trust of the internet.

This decline can be reversed. We can build a trust layer that allows for the building of trust.

The journey towards building a trust layer on the internet is both challenging and essential. As explored throughout this report, the convergence of technical trust and human trust is the cornerstone of achieving trusted interactions.

It's evident that while technology provides the tools (technical trust), governance (human trust) provides the rules to maintain trust.

Trust registries anchor trust, merging technology and human trust. They offer authoritative answers within digital ecosystems, guiding participants on who and what to trust.

The concept of a RoR offers a solution similar to DNS for inventorying and discovering trust registries, creating an interconnected web of trust across ecosystems. This novel approach can further bolster the credibility and utility of trust registries.

While our journey is far from complete, the path ahead is becoming clearer. Technical trust, human trust, and the role of trust registries have been delineated, and we now understand the essential elements required for a deeply trustworthy internet.

Trust is the foundation upon which we will build a future where the internet is once again a place we can rely on, and trusted interactions are the norm, not the exception.

In the end, the success of the emerging trust layer will be measured not only by its technical robustness but also by the trust it enables in our digital relationships. It is our shared responsibility to build and maintain this trust layer for the betterment of digital ecosystems and society as a whole.

Appendix A - Glossary

Anti-Money Laundering/Combating the Financing of Terrorism (AML/CFT) (International Monetary Fund)

Controls, when effectively implemented, that mitigate the adverse effects of criminal economic activity and promote integrity and stability in financial markets.

CAP Dimensions

Acronym referring to three critical aspects of trusted interactions - **C**onfidentiality, **A**uthenticity, and **P**rivacy. These dimensions are vital for assessing and ensuring the trustworthiness of interactions within digital ecosystems.

Decentralized Identifiers (DIDs) (W3C)

A globally unique persistent identifier that does not require a centralized registration authority and is often generated and/or registered cryptographically. The generic format of a DID is defined in 3.1 DID Syntax. A specific DID scheme is defined in a DID method specification. Many—but not all—DID methods make use of distributed ledger technology (DLT) or some other form of decentralized network.

Ecosystem Governance Framework (GHP Glossary)

A governance framework for governing an entire ToIP Layer 4 digital trust ecosystem. An EGF may be either a general EGF or a specific EGF.

Human Trust

The confidence we have in the human interpretation of digital interactions. Human Trust is based on governance, acknowledging that technology alone cannot fully address the complexities of trust in digital interactions.

Interoperability (Wikipedia)

A characteristic of a product or system, whose interfaces are completely understood, to work with other products or systems, at present or in the future, in either implementation or access, without any restrictions.

Know Your Customer (KYC) (Wikipedia)

Guidelines and regulations in financial services require professionals to verify the identity, suitability, and risks involved with maintaining a business relationship with a customer. The procedures fit within the broader scope of anti-money laundering (AML) and combatting the financing of terrorism financing (CFT) regulations.

Registry of Registries (RoR)

A system (centralized or decentralized) that maintains information about various trust registries within a digital ecosystem.

Technical Conformance Suite

A set of tests, criteria, and/or procedures to verify whether a system conforms to the technical standards, protocols, and requirements of an ecosystem.

Technical Trust

The confidence associated with the technical aspects of information exchanged online. Technical Trust is the ability to prove that information originates from a specific source, and has not been tampered with.

Trust Decision (GHP Glossary)

A decision by a party about whether or not it will engage in an interaction or transaction with another party, which includes a determination by the first party whether the risk it runs is acceptable (given its risk appetite).

Trust Layer

The emerging layer of the ToIP Stack that will enable digital interactions to be trusted.

Trust Spanning Layer

In the case of the ToIP Trust Spanning Protocol, the trust spanning layer enables cryptographically verifiable data to flow between any two endpoints regardless of their local trust domain.

Trust Spanning Protocol (TSP)

Protocol to enable universal end-to-end communication among all Endpoint Systems using trusted messages.

Trusted Information

Information that has been evaluated and meets established standards, governance rules, and conformance criteria within an ecosystem.

Trusted Interaction

An interaction becomes trusted when both the data being shared and the relationship between the parties is reliable. This is the ultimate goal, where data exchange becomes a secure and dependable process.

Verifiable Credential (ToIP Core)

A tamper-evident credential whose authorship by an issuer can be cryptographically verified. Verifiable credentials can be used to build verifiable presentations, which can also be cryptographically verified. The claims in a credential can be about different subjects.

Appendix B - Trusted Interactions In Detail

Let's define "trusted interactions" in general terms before delving into specific details. Let's start with how we can elevate information to an authenticated level, which is required for higher levels of trust to be established (i.e. if we can't agree that information came from a source and hasn't been tampered with, we can't build trust on top of that).

We have been able to call information (or its component data) "authentic" for some time. The concept of a Verifiable Credential ("A tamper-evident credential whose authorship by an issuer can be cryptographically verified."[36]) helps establish the authenticity of information. We can verify that the information has not been tampered with and originates from a particular identifier (a DID). That enables:

- **Authentic Information** - information that can be verified as coming from a trusted source and hasn't been altered or tampered with in any way. Imagine it as a data package sealed with a digital mark of authenticity.

Another challenging aspect of interactions is the lack of context. Why are these parties exchanging information? The purpose behind the exchange remains unclear.

The reason for an interaction is crucial. It defines the assurance we need about the other parties and the information being exchanged.

That requires a Reason - the "why" behind our information exchange. There is a **reason** that two (or more) parties are exchanging information.

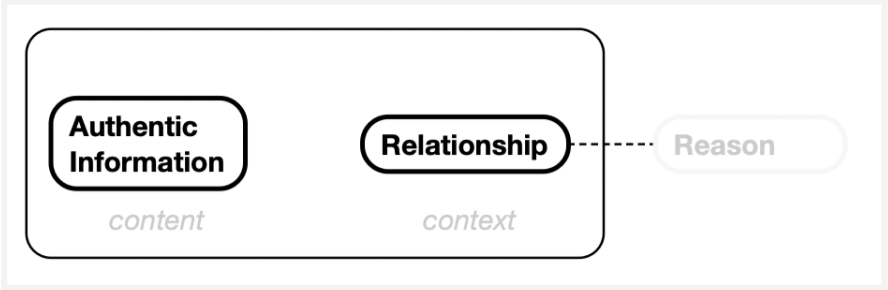

***Fig. B.1** - A Reason for a Digital Interaction*

Further context is added when we look at the governance that applies to the information exchange. The parties need to know that governance (e.g., rules, policies, regulations, legislation) is in place to guide the exchange. This helps manage risks, provide assurances, and establish accountability.

We know the importance of accurately identifying all parties involved in an interaction. The level we need to identify them will vary - from completely anonymous through pseudonymous to fully identified at a very high assurance level.

[36] Trust Over IP Foundation. "verifiable credential." GitHub, 30 August 2022, https://github.com/trustoverip/toip/wiki//verifiable-credential. Accessed 28 September 2023.

When we have the context of the reason that information is being exchanged under a particular governance set, we get:

- **Trusted Relationships** - trust in a relationship is established when we can verify the identity of the parties involved and ensure that they have the authority, reason and governance to engage in the data exchange. Knowing whom we're dealing with is crucial for building trust.
- **Trusted Information** - when we understand a party's authenticity and governance context (e.g., a particular party is/isn't authorized to perform a particular action), we can establish a higher level of trust with the authentic information they produce, hold, or verify.

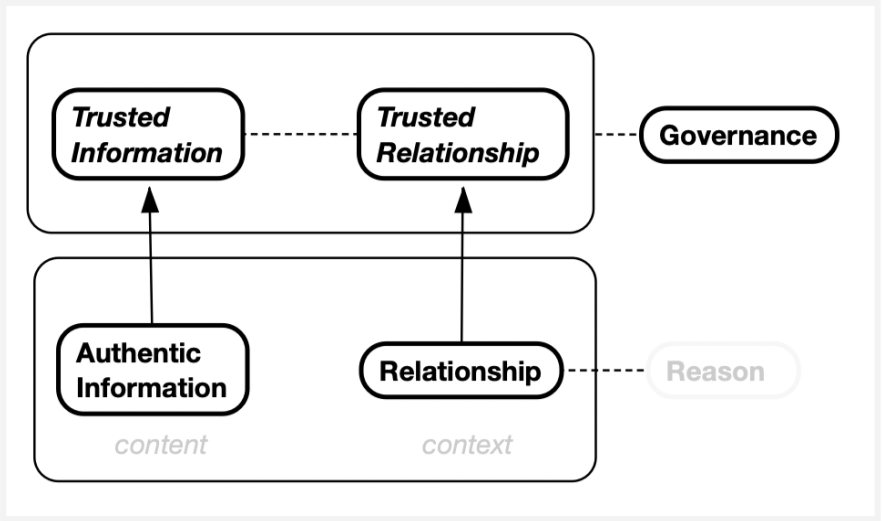

***Fig. B.2** - Transformation of Information and Relationships through Governance*

Governance transforms the interaction.

- Authentic Information can become Trusted Information; and
- A Relationship can become a Trusted Relationship.

This leads to:

- **Trusted Interactions** - A digital interaction where Trusted Information is exchanged in the context of a Trusted Relationship.

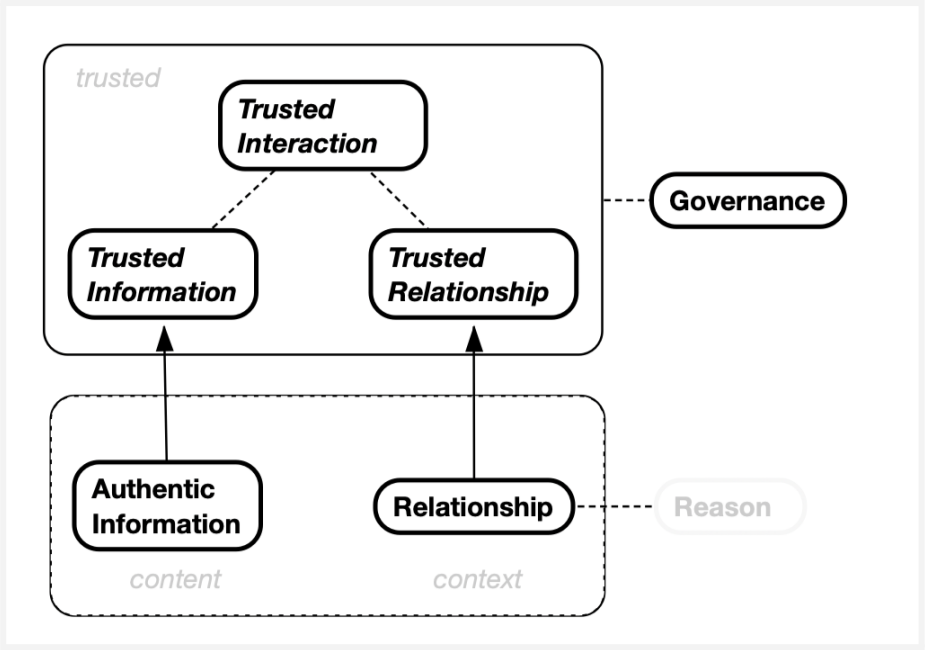

***Fig. B.3** - Trusted Interaction*

We can shift to a specific type of information exchange - where Trusted Information is exchanged in the context of a Trusted Relationship - under a particular Governance scheme:

- **Trusted Interaction** - An interaction becomes trusted when both the data being shared, and the relationship between the parties is reliable. This is the ultimate goal, where data exchange becomes a secure and dependable process.

By following these principles, businesses can ensure that their data exchanges are conducted with integrity, security, and transparency.

Trust is the bedrock of successful partnerships, and by embracing these concepts, businesses can build strong and reliable relationships with their partners, customers, and stakeholders. This, in turn, fosters growth, collaboration, and confidence in the ever-changing landscape of data-driven business.

Figure B.4 below depicts how this concept of moving from Authentic Information and Relationships can transform into fully trusted interactions.

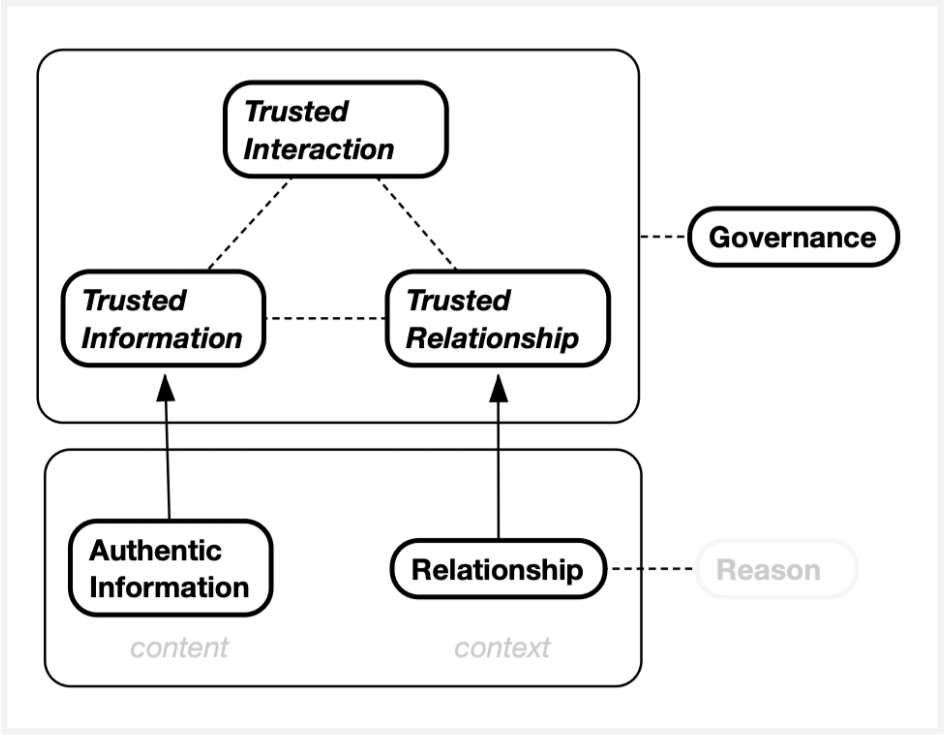

Fig. B.4 *- Evolution to Trusted Interactions*

To fully understand these trusted interactions, we need to understand what we mean by this loaded term “trust.”

B.1. Trusted Interactions and Trust Registries

In our trusted interactions, the parties involved want to build sufficient trust that the other party is appropriate to interact with. This may mean identifying them sufficiently or getting assurances that they are authorized to perform their part of an interaction.

Sometimes we need to identify each other, but often we don’t. We may just need assurance that the other party is authorized for our specific purpose. Whether we need to identify a party or just know they are appropriately authorized, we want to ensure that they have Authenticity.

Additionally, we need to know how to connect to establish a trusted interaction. Chatting, for example, may require minimal setup. Connecting to sign a contract with many parties and government approval could require an incredible amount of information just to connect the systems.

We want to know that the trusted interaction meets the governance requirements.

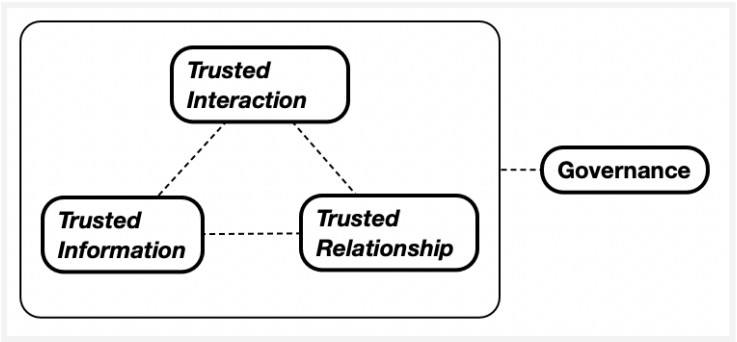

Fig. B.5 - *Establishing Trusted Interactions*

B.2. Trusted Relationships and Trust Registries

Once we have a place where we can ask questions and reasonably “trust,” we are at the point of being able to trust data and relationships.

Knowing that one or more trust registries recognize the issuer of a government credential is powerful. It’s hard for fraudsters to get registered in official trust registries.

When we look at the Trusted Relationship in a trusted interaction, we need to know if the parties in that relationship are authentic. Are they who they say they are, and do they have the authorizations required for a particular interaction?

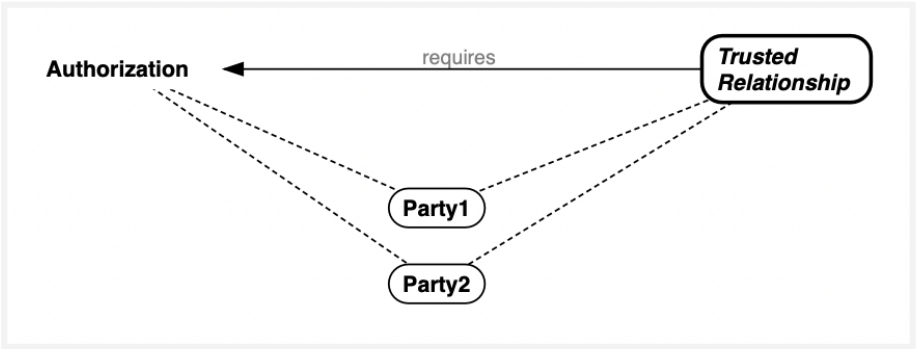

Fig. B.6 - *Authenticity and Authorization in Trusted Relationships*

Appendix C - CAP and Trust - Human & Technical

Certain inherent cryptography and internet capabilities bolster the technical trust. We know that cryptography allows relatively high levels of confidentiality, the first piece of the CAP. Human trust will come into play depending on the requirements for authenticity and privacy.

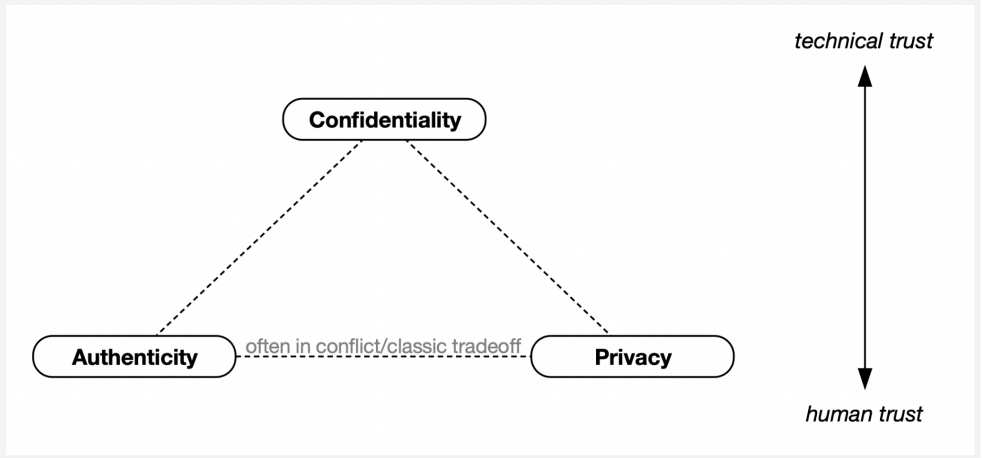

Fig. C.1 - *Balancing Confidentiality, Authenticity, and Privacy*

C.1. Confidentiality - Encryption and Messaging

Confidentiality involves ensuring that unauthorized parties cannot access the information being exchanged. Cryptography handles the confidentiality of most of the information in a digital interaction. However, it does not provide the parties' confidentiality, as their traffic can't be hidden.

Appropriate messaging approaches can provide the confidentiality of the parties involved in a digital interaction.

As stated elsewhere, Confidentiality can be assured mostly on a technical trust level.

C.1.1. Confidentiality - Technical Trust

Confidentiality is well handled on a mostly technical trust basis via well-known cryptography and messaging. The approach recommended by ToIP involves an end-to-end confidentiality approach. End-to-end confidentiality is well described in a recent ToIP Foundation document[37]:

> *Almost all modern security protocols based on public/private key cryptography use some combination of message signing (for authenticity) and message encryption (for confidentiality). A longstanding question has been: precisely what combination of these two properties produces the strongest security?*

[37] Trust Over IP Foundation. "Mid-Year Progress Report on the ToIP Trust Spanning Protocol." *trustoverip.org*, 9 March 2019, https://trustoverip.org/blog/2023/08/31/mid-year-progress-report-on-the-toip-trust-spanning-protocol. Accessed 7 November 2023

Our second pillar is a firm answer to that question: the signing and encryption pattern that provides the strongest protection against both key compromise impersonation (KCI) and sender impersonation of the ciphertext is called ESSR (for Encrypt Sender's key then Sign Receiver's key). ESSR was first defined in a 2001 paper by Jee Hea An and is well explained in these three Neil Madden blog posts about public key authenticated encryption: PKAE1, PKAE2, PKAE3.

The bottom line: by binding the sender's public key inside the encrypted ciphertext and binding the receiver's public key in the enclosing signed plain text, an adversary is prevented from forging messages that compromise either authenticity or confidentiality. So the trust spanning protocol can achieve both strong authenticity and strong confidentiality by applying ESSR to all messages that require both properties.

Implementing end-to-end confidentiality allows the highest level of technical trust, and when combined with human trust, it can further enhance overall trust.

C.1.2. Confidentiality - Human Trust

As mentioned in the technical trust discussion above, Confidentiality is largely handled on a technical level. While there are exceptions to this technical focus, there is one exception[38] of note:

- Governance may impose limitations on the acceptable cryptographic suites. For example, many US and Canadian government systems require that encryption modules comply with the Cryptographic Module Validation Program[39].

C.2. Authenticity - Identifying the Parties

Authenticity encompasses various dimensions that differ when examined from both a technical trust and a human trust perspective:

- Technically, we can know the identifier of an endpoint and that the identifier is being used to encrypt and sign messages. This means we can technically trust the identifier.
- On a purely technical basis, we can't know the IDENTITY of an endpoint. That requires human trust. It correlates the identifier (technical trust) with the Party we interact with (human trust) using some kind of governance.

The technical approaches above increase the level of trust associated with a digital interaction.

The ToIP mid-year update referenced earlier states, "*an adversary is prevented from forging messages that compromise either authenticity or confidentiality*[40]." This definition of authenticity is technical, applies well to the technical trust requirements, and allows systems to maintain authenticity on that technical trust level.

[38] There will certainly be other important aspects for human trust, but this one stands out to the author.
[39] Computer Security Resource Centre. "Cryptographic Module Validation Program." *csrc.nist.gov*, 2016, https://csrc.nist.gov/projects/cryptographic-module-validation-program. Accessed 7 November 2023.
[40] Trust Over IP Foundation. "Mid-Year Progress Report on the ToIP Trust Spanning Protocol." *trustoverip.org*, 9 March 2019, https://trustoverip.org/blog/2023/08/31/mid-year-progress-report-on-the-toip-trust-spanning-protocol. Accessed 7 November 2023.

However, it requires human trust to be layered on top of the technical trust to ensure that the highest levels of trust can be established. More on that below.

C.2.1. Authenticity - Technical Trust

The general approach of using decentralized identifiers allows parties to meet some technical authenticity aspects. In particular:

- Decentralized Identifiers (DIDs) - the use of DIDs, which is an approved W3C standard[41], allows a party to control their identifier and associated keys. Additional benefits can be established using an autonomous identifier, a specialized kind of DID.
- Proving Control - Using the keys associated with the DID proves that an identifier is controlled by the party involved in an interaction.
- Using Authenticated Communications - sharing information across a confidential channel to all parties allows continuous authentication. Various well-understood approaches already use this approach (e.g. DIDComm[42], Key Event Receipt Infrastructure (KERI)[43], Message Level Security[44] at IETF).

C.2.2. Authenticity - Human Trust

Authenticity is about knowing that the party you are interacting with is who they say they are and appropriate for your interaction.

- Authorization - The use of approaches that provide evidence that a party has the necessary authority to carry out the requested interactions. There are two key approaches to this:
    - Trust registries - querying an authoritative source about a DID's authority is a straightforward pattern. This approach provides an unambiguous signal that can be layered with other queries (e.g. querying multiple trust registries).
    - A credential may be presented by the Party at the other end that proves their bona fides sufficiently for an interaction. This approach adds a layer of complexity as the issuer's (of the credential) authority must also be confirmed.
- Authentication - we need to know that the Party we are communicating with is the appropriate party. While we can prove on a technical basis that a party is controlling an identifier - do we know that the party is who we think they are? This can't happen by technical means unless we weave governance into the process and have systems of record that we can query. A trust registry helps here and may be sufficient to create the human trust element.

---

[41] World Wide Web Consortium. "Decentralized Identifiers (DIDs) v1.0." *w3.org*, 2022, https://www.w3.org/TR/did-core. Accessed 7 November 2023.
[42] Decentralized Identity Foundation. "DIDComm Messaging v2.1." *identity.foundation*, 9 March 2019, https://identity.foundation/didcomm-messaging/spec/v2.1. Accessed 7 November 2023.
[43] Smith, Samuel M. "Key Event Receipt Infrastructure (KERI)." *arxiv.org*, 3 July 2019, https://arxiv.org/abs/1907.02143. Accessed 7 November 2023.
[44] Sullivan, Nick, and Sean Turner. "Messaging Layer Security: Secure and Usable End-to-End Encryption." *ietf.org*, 29 March 2023, https://www.ietf.org/blog/mls-secure-and-usable-end-to-end-encryption/. Accessed 28 September 2023.

C.3. Privacy - Following The Rules (and Proving It)

The confidentiality and authenticity aspects of CAP lean heavily towards technical trust. They allow for a certain level of Privacy, but privacy is a much more governance-related (i.e. human trust) concept.

While technical approaches to privacy exist, they are principally in place to provide statements of compliance and consent. They do not create privacy.

C.3.1. Privacy - Technical Trust

For this report, we won't dwell too much on the technical aspects that apply to privacy, as privacy is evolving and fairly early in full technical enablement.

While many techniques exist to enforce privacy, the fundamental way for Privacy happens in the human trust element. It is the governance that imposes and implements approaches to privacy.

C.3.2. Privacy - Human Trust

While we can technically enhance privacy, fully controlling the release of exchanged information through technology remains impossible. The human trust aspects here are governance (formal and informal) related. Privacy is handled by legislation (e.g. privacy laws), regulation, governance/trust frameworks, and social mores.

In situations where Privacy is concerned, a formal governance framework may be the best approach. Parties to these frameworks need to understand the expectations for privacy and the penalties for violating the privacy agreements.