ayra

trusted relationship network

The Ecosystem-of-Ecosystems Model for Decentralised Digital Trust Infrastructure

A White Paper from the Ayra Association

https://ayra.forum/

Version 3.0
16 January 2025

Darrell O'Donnell
Drummond Reed
Wendy Seltzer
Andor Kesselman
Dave Poltorak

Audience and Purpose

This white paper is written for Internet professionals of all kinds—CIOs and CISOs; security, privacy, and trust architects; policy experts; product managers; and developers—who want to quickly understand the new layer of decentralized digital trust infrastructure Ayra (eye-rah) is building.

Abstract

Digital wallets and cryptographically-verifiable digital credentials represent powerful new tools for establishing digital trust relationships. However they are already facing the problem of how a credential issued in one digital trust ecosystem can be accepted in a different one.

Tackling this problem in a decentralized architecture that empowers every ecosystem to self-govern, yet still enable cross-ecosystem interoperability, requires a new layer of digital public infrastructure. This layer resembles DNS—and needs similar global multi-stakeholder governance—but also differs from DNS in several crucial ways.

The paper explains why this layer needs an "ecosystem of ecosystems" model to reflect the reality that many trust relationships are non-hierarchial. It then proposes that the steps to build and deploy this layer directly parallel those required to stand up DNS 50 years ago. This includes the critical role of forming and deploying the first **trust clusters**—a role the Ayra Association is designed to fill as a neutral international Swiss nonprofit association.

Recommended Pre-Reading

This short pre-read quickly explains the overall Ayra problem and solution space:

Introduction to Ayra (7 pages): ➤ Why has digital trust eroded so badly? ➤ Why is this such a threat to AI? ➤ How the "friend-of-a-friend" pattern works in the real world; ➤ How credit card networks solved the acceptance problem; ➤ How Ayra is and is not like a credit card network ➤ What is a trust cluster? ➤ What are the first trust clusters that are forming? ➤ How do I get started building the trust cluster I need?

Recommended follow-up reading:

The Special Network Effects of Ayra Network Credentials (21 pages): ➤ The seven design principles for Ayra network credentials; ➤ The special roles of the First Person and Org ID credentials, ➤ Why Ayra network credentials have five compounding network effects.

Ayra Technical White Paper (50 pages): ➤ Core technical architecture of the Ayra Trust Network; ➤ The ToIP Trust Registry Query Protocol (TRQP); ➤ The Ayra Conformance Test Suite (CTS); ➤ The Ayra Type Catalogue; ➤ The Ayra Technical Working Group.

Table of Contents

Introduction **5**
Terminology **5**
PART ONE: The Interoperability Challenges of Digital Trust Ecosystems **6**
What is a digital trust ecosystem? 6
The challenge of verifying digital credentials 6
Verifiable identifiers 8
Trust registries (aka trust lists or trusted lists) 8
Credential status registries 10
The challenge of interoperable trust registries 10
PART TWO: The Ecosystems of Ecosystems Model **13**
Comparison with DNS 13
Hierarchy vs. heterarchy 14
Resource discovery vs. trust verification 15
The Ayra Trust Network 16
Step One: Software 16
Step Two: Governance 16
Step Three: Recognized Ecosystems (Trust Clusters) 17
PART THREE: The Trust Cluster Formation Process **18**
Stage 1: Informal Cluster Discussions 18
Stage 2: Cluster Projects 18
Stage 3: Cluster Metaregistries 19
Stage 4: Cluster Governance Frameworks 19

Acknowledgements

The authors would like to thank the many members of the Ayra Association who contributed to this paper, and in particular representatives of the Member Advisory Council and the Ayra Ambassadors. This represents a collective effort by the entire community.

Introduction

National and international digital wallet and digital credential projects such as the European Digital Identity Wallets initiative, the Bhutan National Digital Identity program, the Province of British Columbia Digital Trust program, and Findynet are proof that decentralised digital trust infrastructure is here to stay.

These initiatives also underscore how critical it is becoming for these digital wallets and credentials to interoperate at Internet scale—the same way Internet browsers and web servers interoperate everywhere, or that email clients and email servers interoperate everywhere. Addressing this need is the reason for the recent announcement of the Ayra Association and the Ayra Trust Network.

As stated on the Ayra website:

> *The purpose of the Ayra Trust Network is to drive the growth of safe, secure, interoperable, and sustainable digital trust ecosystems that connect and enrich our world.*

One of the core premises of the Ayra Trust Network is that, in order to be globally scalable and decentralised as whole (even though individual ecosystems may be centralised or federated), the overall design needs to follow an **ecosystem-of-ecosystems model**. The purpose of this paper is to explain that model and how the Ayra Trust Network is pursuing its development and implementation by facilitating formation of the initial **trust clusters**.

Terminology

Most of the terms used in this paper—and more broadly across decentralised digital trust infrastructure—are defined in the ToIP Glossary maintained by the Trust Over IP (ToIP) Foundation. Terms specific to Ayra infrastructure are defined inline.

PART ONE: The Interoperability Challenges of Digital Trust Ecosystems

What is a digital trust ecosystem?

The term "ecosystem" has become widely used in digital trust infrastructure because it describes how a variety of actors (including competitors) can organically cooperate to achieve a mutual set of digital trust objectives. Here is the formal definition of the term **digital trust ecosystem** in the ToIP Glossary:

> *A digital ecosystem in which the participants are one or more interoperating trust communities. Governance of the various roles of governed parties within a digital trust ecosystem (e.g., issuers, holders, verifiers, certification bodies, auditors) is typically managed by a governing body using a governance framework as recommended in the ToIP Governance Stack. Many digital trust ecosystems will also maintain one or more trust lists and/or trust registries.*

When applied to digital wallets, an ecosystem consists of the issuers, holders, and verifiers (a.k.a. relying parties) for a specific type of digital credential (e.g., drivers license, proof of insurance, employee ID, loyalty card) plus the governing body (formal or informal) for those collaborating parties. This is illustrated in figure 1, where the upper half is often referred to as the "trust triangle" and the full diagram as the "trust diamond" or "governance diamond".

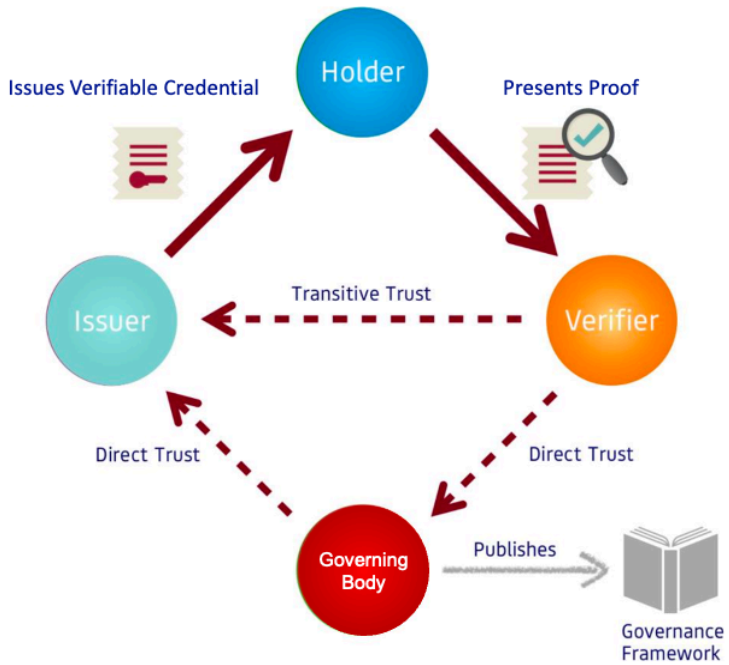

Figure 1: The core roles in a credential-based digital trust ecosystem.

The challenge of verifying digital credentials

With the paper or plastic credentials we carry in our physical wallets, verification is typically performed by a human agent who manually inspects the credential. One exception is when a physical credential contains machine-readable information, such as a barcode or encoded biometric. In that case, verification can be at least partially automated—typically by a dedicated system or network designed to work with a specific type of credential (for example, a U.S. TSA airport security system verifying a U.S. driving licence or passport).

When this same verification task moves into the digital sphere, it must be performed by software—almost always without a human in the loop. This forces us to spell out the steps that actually take place in the process of human verification. The process begins with a software agent requesting presentation of a credential from a holder's digital wallet.

The very first step is determining whether the party to whom the credential was issued is the party actually presenting it for verification—a step often called holder binding. There are multiple approaches to holder binding depending on the level of assurance required—they can include verifying a biometric bound to the credential, verifying a decentralized identifier (DID) bound to the credential, or requesting proof-of-presence from a third party provider.

After holder binding, the software agent must perform the four verification tasks in table 1:

#	Task	Description	Required Support
1	**Verify digital signature**	Verify that the credential presentation has a valid digital signature produced by the private key bound to the public key.	None (can be verified using cryptography alone)
2	**Verify public key**	Verify that the public key is controlled by the credential issuer.	Verifiable identifiers (e.g., DIDs)
3	**Verify issuer (or verifier) authorization**	Verify that the issuer (or verifier) is authorised within an ecosystem to issue (or verify) this specific type of digital credential.	Trust registry (aka trust list)
4	**Verify credential status**	Verify that the credential is not suspended, expired, or revoked.	Revocation registry

Table 1: The four tasks required to verify a digital credential

The first of these four tasks the software agent can perform using cryptography alone. The other three generally require some type of supporting system as defined in this section. Because these supporting systems are critical to trust establishment, they must be either operated by—or authorised by—the authoritative source or governing body for an ecosystem.

Verifiable identifiers

Verifiable identifiers (VIDs) are the "atomic building block" of decentralised digital trust infrastructure. Here is the definition of verifiable identifier in the ToIP Glossary:[1]

> *An identifier over which the controller can provide cryptographic proof of control. Each type of VID defines a specific means for discovering the public key, network endpoints, or other metadata necessary to prove control. Decentralized identifiers (DIDs) are a W3C standard for VIDs. VIDs are the cryptographically verifiable identifiers used in the ToIP stack.*

In short, to perform task #2 in table 1, the software agent needs to perform the steps required to verify the specific type of verifiable identifier being used by the issuer. This includes verifying the associated public key(s). Generally speaking, *this is not an ecosystem-specific verification task*. The market is expected to support a relatively small number of VID types (e.g., DID methods) that will work across many different ecosystems.[2]

Trust registries (aka trust lists or trusted lists)

The third task in table 1 requires the verification agent to check with an authoritative source to determine if the issuer is authorised to perform that role in the context of a specific ecosystem. Answering such queries is the job of a trust registry as shown in Figure 2.

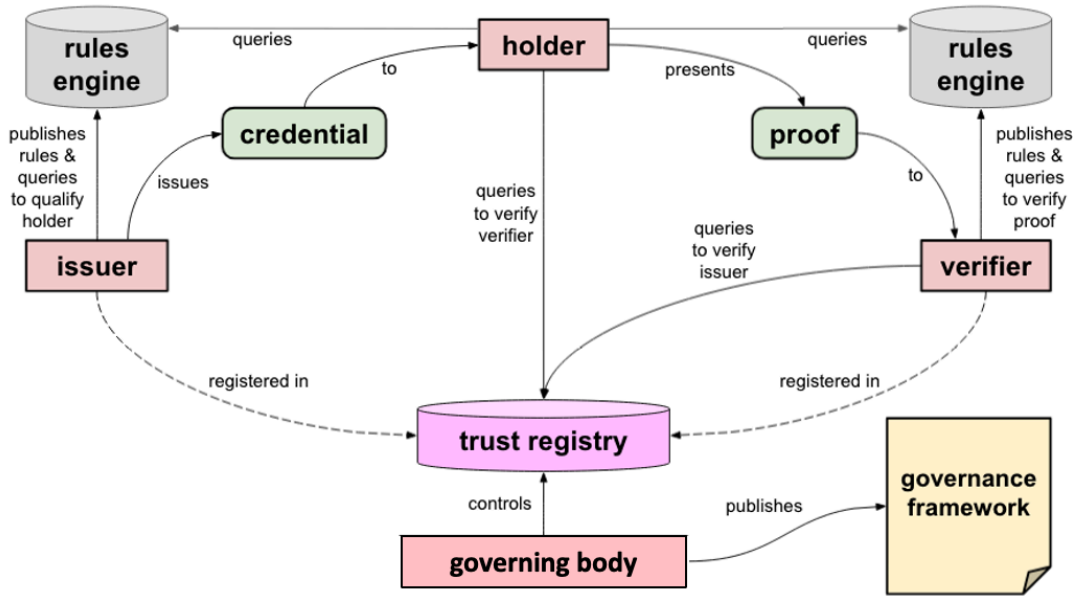

Figure 2: The role of a trust registry (aka trust list or trusted list) in a digital credential ecosystem

[1] The W3C Decentralized Identifiers (DIDs) 1.0 specification and the W3C Verifiable Credentials Data Model 1.1 specification both use the term verifiable data registry to describe the supporting system used to verify a verifiable identifier . However, this is somewhat of a misnomer, because verification may not involve a registry at all, particularly for self-certifying identifiers (SCIDs).

[2] Some ecosystems may have specific requirements for the VID types or DID methods they will accept.

In the EU, trust registries are known as **trusted lists**, and the EU publishes them in the List of Trusted Lists (LOTL) as shown in Figure 3. In other PKI infrastructure, trust registries are also known as **trust lists**.

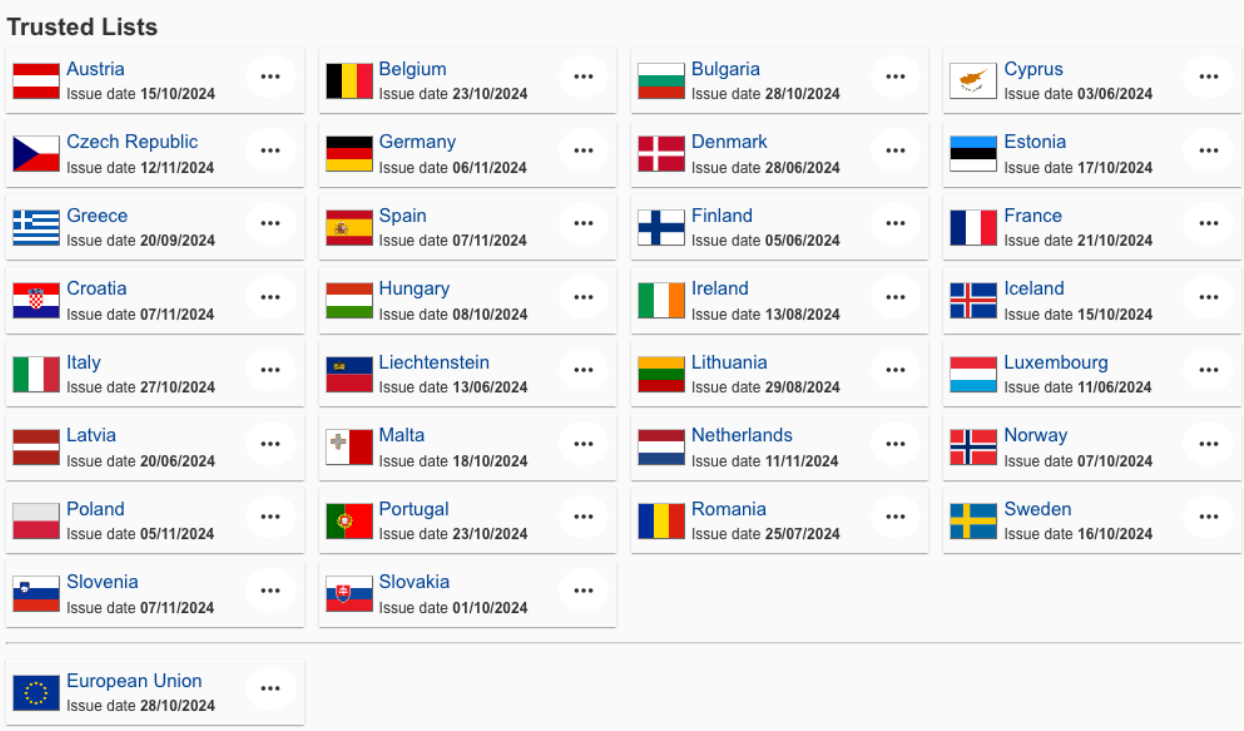

Figure 3: Example from the EU List of Trusted Lists (LOTL)

In contrast with verifiable identifiers, whose method of verification is generally independent of any specific ecosystem, trust registries are always ecosystem-specific because their purpose is to answer queries for authorization data in the following general form:

> Does entity X have authorization Y in the context of ecosystem Z?

Following are common examples of these types of trust registry queries:

- Is hospital X authorised to issue health credential Y in ecosystem Z?
- Is company X authorised to verify employment credential Y in ecosystem Z?
- Is auditor X authorised to conduct security audit Y in ecosystem Z?
- Is certification body X accredited to certify Y in ecosystem Z?
- Is ecosystem X recognized as a trust registry by ecosystem Z?

To perform this verification task across different ecosystems, a verification agent must be able to ask these verification questions using an **open standard trust registry query protocol**. This interoperability requirement will be further discussed below.

Credential status registries

The final verification task in table 1 is checking the credential status. Depending on the credential and the governance framework, status options might include:

- **Active**—the credential is currently valid.
- **Expired**—the credential is not valid past a specific date.
- **Suspended**—the credential is valid but the authorization it grants is temporarily revoked (such as might apply to a driving licence or professional certification).
- **Revoked**—the credential is no longer valid. This status could apply to any credential—even those with expiration dates.

Credential status registries are essentially specialised types of trust registries because they are designed to answer one specific type of verification query (task #4 in table 1):

> What is the status of credential X?

The key challenge is how to answer this question without putting a globally unique identifier in the credential ("X" in the example above) because that can serve as a global tracking beacon for that credential. The industry is actively working on revocation registries that solve this problem using privacy-preserving technologies such as zero-knowledge proofs.[3]

For the purposes of this paper, we will consider revocation registries to be a specialised type of trust registry.

The challenge of interoperable trust registries

When a trust registry is needed for only a few specific types of credentials within a single ecosystem—such as the U.S. TSA example given earlier—there is little need for interoperability. However, when thousands of different types of digital credentials become available from millions of different ecosystems around the world, the need for interoperability between trust registries becomes overwhelming.

A common example is educational credentials. Many employers around the world would like to be able to easily verify the degrees, certificates, or other educational credentials of a prospective employee or contractor. But how can an employer be expected to know all the accredited educational institutions in the world? Or even in one country?

The need to verify authorizations and accreditations across ecosystems is the same overwhelming need for interoperability that produced the Internet fifty years ago. We had thousands of local area networks, and we needed to connect them so data could flow between any two devices anywhere in the world.

[3] See the Cryptographer's Response on the EU's Digital Identity Architecture Reference Framework (ARF).

Now we need to do the same thing to build the **Internet of Trust**. We will soon have thousands of digital trust ecosystems around the world, each issuing digital credentials to millions of people and businesses (even to devices and bots). Can we agree on a standard, secure, highly scalable way to interconnect trust registries so that verifying digital credentials on the Internet becomes as easy and fast as resolving a domain name?

The technical half of this challenge–while not trivial—is the easy part. Work on open standard trust registry protocols has been ongoing since the global pandemic created an urgent worldwide demand for interoperable COVID credentials. That led to the founding of the Trust Registry Task Force (TRTF) at Trust Over IP (ToIP). Four years later, the TRTF has produced the second generation of its Trust Registry Query Protocol (TRQP) that can standardise queries across many different types of trust registries as illustrated in Figure 4.[4]

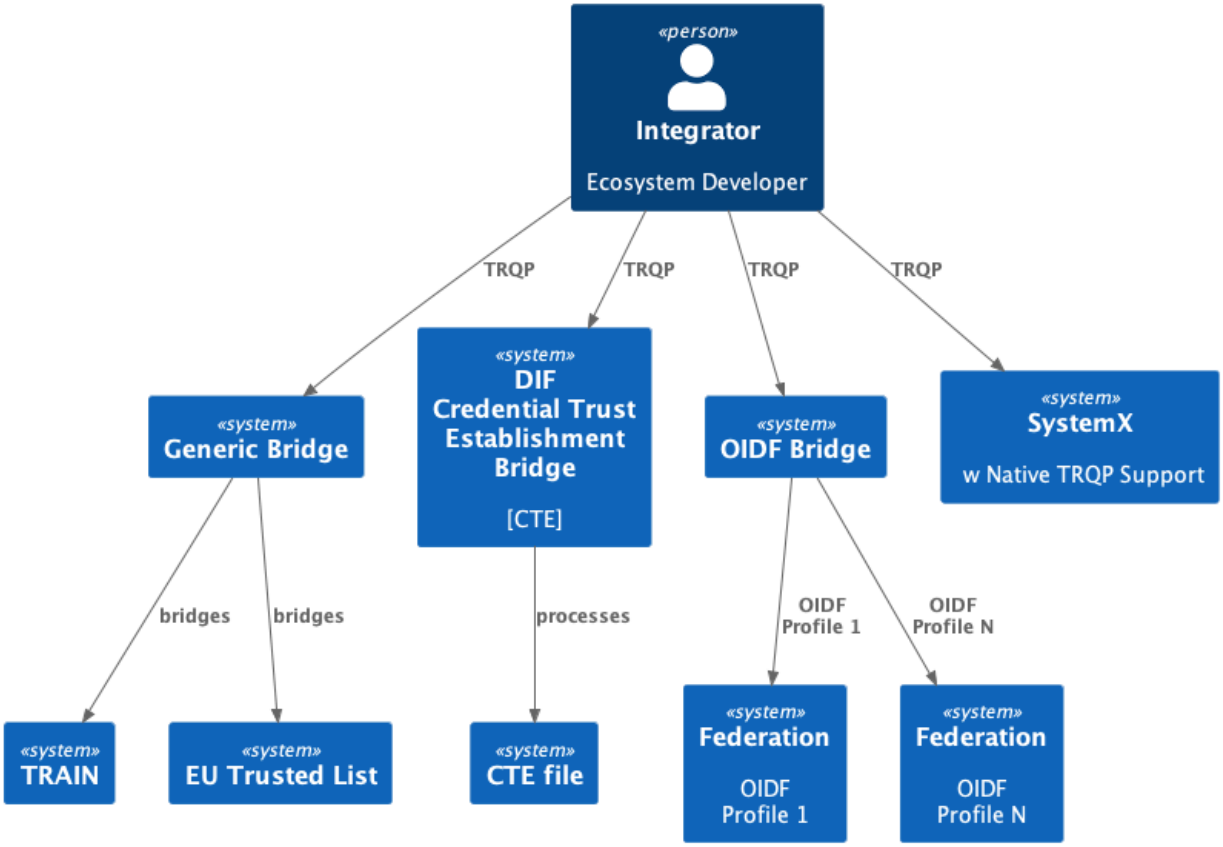

Figure 4: The TRQP standardises access to trust registry data from many different sources.

The harder part of the problem is in fact the *human part*. In other words, while bridging local area data networks together to form the Internet was 90% technology and 10% governance, bridging digital trust ecosystems is just the opposite.

[4] Source: Annex C of the TRQP specification. For more detailed information on the technical architecture of trust registries and metaregistries in the Ayra Trust Network, see the Ayra Technical White Paper.

This was the reason ToIP was founded in 2020. Although the left side of the ToIP stack shown in figure 5—the technology stack—is patterned directly after the Internet's four-layer TCP/IP stack, the right side—the governance stack—shows how technology must be bound to human trust in the form of explicit governance frameworks to solve the problem of interoperability among digital trust ecosystems.

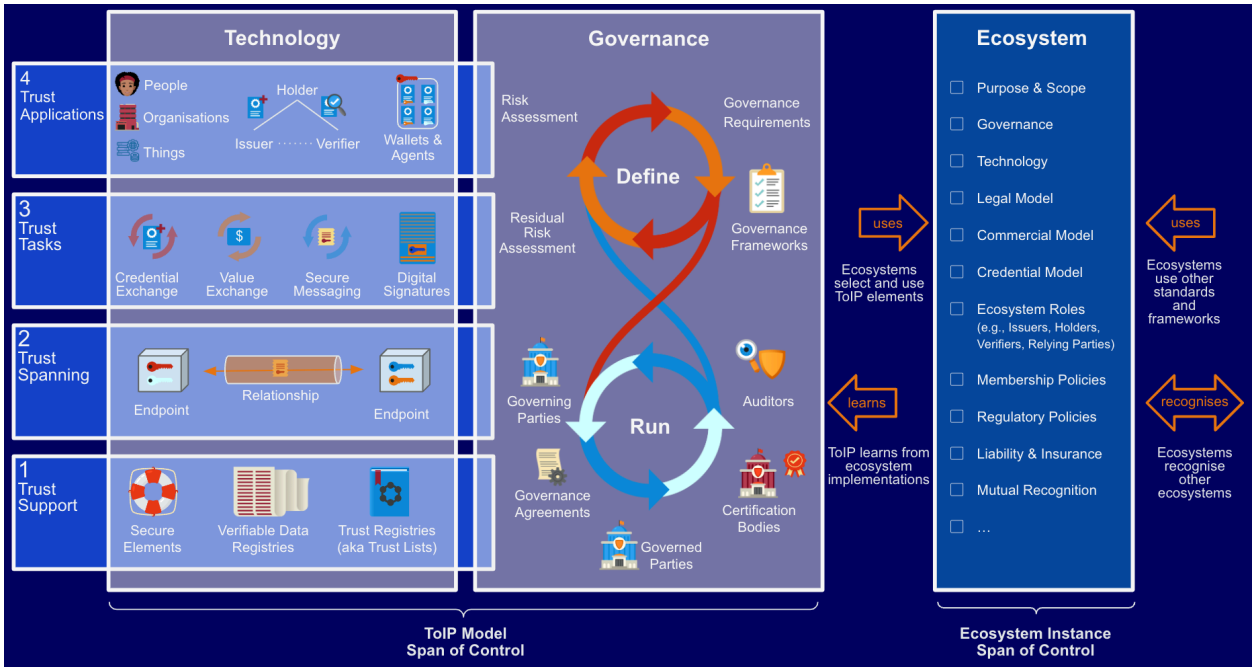

Figure 5: Interoperable digital trust ecosystems require both technology and governance.

Only when technology and governance are brought together in a consistent model can we finally tackle the challenge of Internet-scale interoperability between digital trust ecosystems. Which leads us to the next question: what does that scalable infrastructure look like?

PART TWO: The Ecosystems of Ecosystems Model

Comparison with DNS

When the domain name system (DNS) was developed to standardise mapping human-friendly names to Internet addresses, its answer to scaling was federation. Rather than one giant centralised list of domain names for the entire Internet, each domain controller maintains its own list of unique names. Each name can point either to a resource inside that domain, or delegate to another domain. This way domain names can be chained together into a "name path".

Figure 4 shows an example for the name path **www.example.org**. Because DNS name paths typically read right-to-left, it shows how a DNS client resolver might first query the ICANN root server for the domain named "**org**". The ICANN root server delegates to the authoritative name server for "org". This pattern continues: the "org" server is queried for the domain named "example". The "org" server delegates to the "example" server. Once the query has reached the authoritative server for the fully qualified domain name, that name server is able to answer questions about the records it manages. In figure 6, this is the final response from the "www" subdomain containing the IP address for that server (by convention a web server).

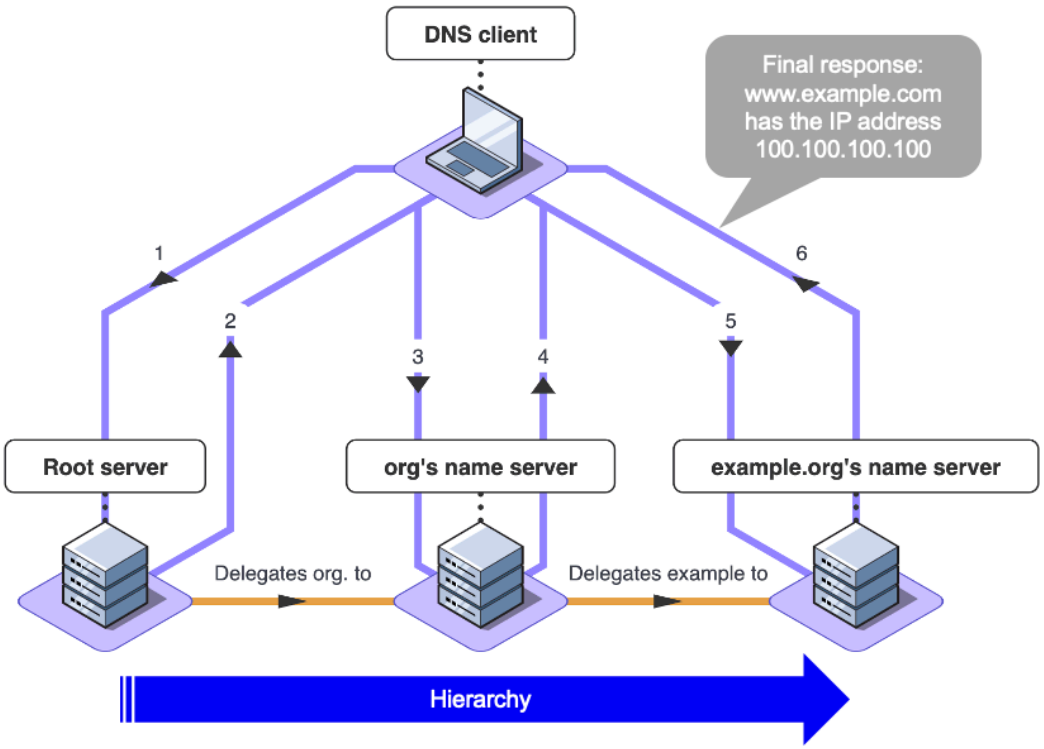

Figure 6: How DNS scales by enabling domain name servers to delegate along a hierarchical name path.

This federation architecture enables each domain controller to manage their own internal naming for their own set of Internet-accessible resources, while also giving them the power to delegate to other domains in order to make them Internet accessible—without requiring the delegating domain controller to manage the names in those delegated domains.

Although not perfect—particularly with regard to security—the DNS has successfully scaled to serve the entire Internet, with an average domain name lookup time in milliseconds.

The same scalability principle can be applied to the trust registries for digital trust ecosystems. We call this the **ecosystem-of-ecosystems** model. However there are several important differences between the ecosystem-of-ecosystems model and the DNS model.

Hierarchy vs. heterarchy

First, the ecosystem-of-ecosystems model is not hierarchical, but heterarchical. In other words, rather than a strict tree like DNS, where each higher level domain delegates to the next domain lower in the tree, in the ecosystem-of-ecosystems model, each ecosystem's trust registry can choose to *recognize* another ecosystem's trust registry in a peer-to-peer fashion.

This recognition can be *unidirectional*, i.e., ecosystem A trusts ecosystem B, but ecosystem B does not trust ecosystem A; or *bidirectional* (aka *mutual*), i.e., A trusts B and B trusts A. Either way, the difference with DNS is that domain name paths traverse *hierarchically down* a domain name tree, while trust chains may also traverse *laterally across* ecosystems—with each being a peer of the others. Figure 7 is a way of illustrating this—the blue arrows indicate heterarchical recognition; the red arrows indicate hierarchical delegation.

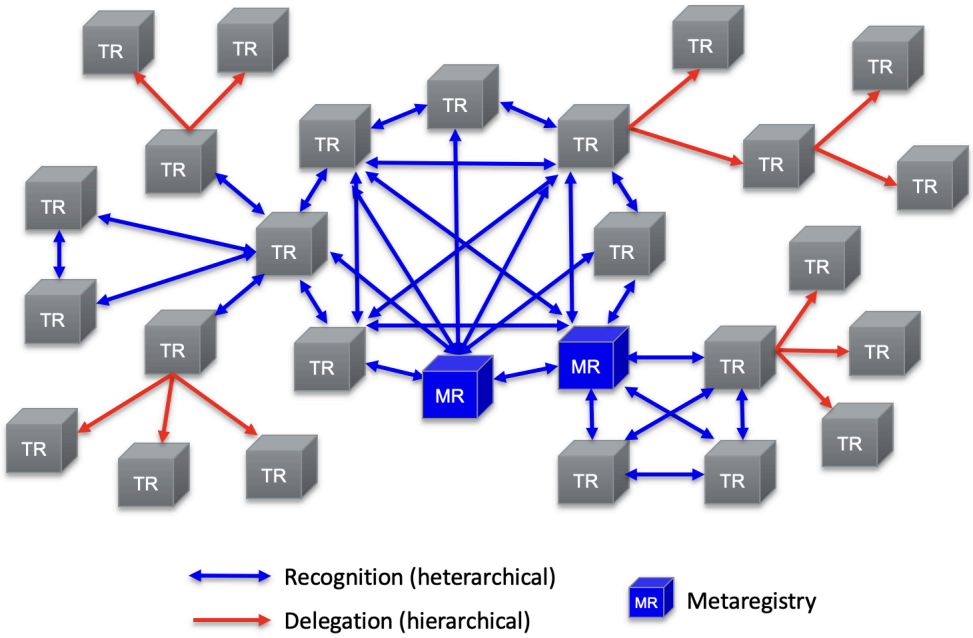

Figure 7: The heterarchical structure of an ecosystem-of-ecosystems model for trust registries.

Users of ActvityPub will recognize this model universe where each ActivityPub server admin makes the decision about recognizing another server ("Will I federate with you?"). While such recognition is typically pairwise, the decision is ultimately up to each server admin.

Of course, hierarchy is a subset of heterarchy, i.e., hierarchical delegation chains can be set up within a heterarchical network. So, for example, the trust registry for a national ecosystem could delegate to the trust registry for a member state ecosystem which could delegate to the trust registry for a local ecosystem. But with the ecosystem-of-ecosystems model, recognition is as viable as delegation, i.e., an ecosystem can recognize the sovereignty of another ecosystem without having any authority over the recognized ecosystem.

Such recognition has the same upsides and downsides as diplomatic recognition in the real world. On the positive side, it enables trust to "flow" over new trust relationships between the two ecosystems. On the negative side, there is always the risk that a recognized ecosystem might change its policies or stop enforcing them. Just as with diplomatic recognition, such risks can be mitigated by the threat of withdrawing recognition and/or by explicit contractual agreements between the ecosystems.

Heterarchies can of course be less efficient than hierarchies for machines to navigate; however modern network searching and indexing technologies should be able to minimise any such inefficiencies. In addition, in an ecosystem-of-ecosystems architecture, efficient navigation can be aided using **metaregistries**. A metaregistry is a trust registry whose primary function is to recognize other trust registries within a particular ecosystem. In figure 7, the two trust registries shown in blue are metaregistries—the only records they contain are pointers to other recognized trust registries within their particular ecosystem-of-ecosystems.

Resource discovery vs. trust verification

The second major difference between the ecosystem-of-ecosystems model and the DNS model is that domain names serve almost exclusively for digital resource discovery, i.e., as pointers to locate the IP address (or other attributes) of some type of Internet service endpoint (e.g., web server, email server, file server, or another name server). *They do not convey trust in a resource* other than signalling that the resource is controlled by the domain name authority—a signal that carries greater strength if the domain uses DNSSEC to digitally sign its name records.[5]

By contrast, the records in a trust registry *convey explicit trust decisions* by the governing body about which ecosystem participants are authorised to perform which roles—and what other ecosystems are recognized as sufficiently trustworthy for their trust registries to be recognized as peers. This is the verification data we need to build the Internet of Trust.

[5] Certain DNS top-level domains (TLDs), such as those allocated to countries, may signal some additional trust depending on the domain controller and the policies governing that TLD.

The Ayra Trust Network

Although the technical protocols for the DNS were codified in the Internet Engineering Task Force (IETF), that alone did not create a functioning global DNS network. Doing so required three additional steps:[6]

1. Development of the open source and commercial software for name servers and resolvers that spoke the DNS protocols (such as the nearly ubiquitous BIND);
2. Establishing a governing body for the DNS top-level domains (TLDs). This eventually became the Internet Corporation for Assigned Names and Numbers (ICANN).
3. Launching an initial set of TLD registries including the two-letter country-code registries (such as .us, .fr, .ca, .jp, .nl) operated by governments (or their delegates) and generic TLD registries (such as .com, .net, .org) operated by various private organizations.

These same three steps are necessary to establish a global network of trust registries.

Step One: Software

The first step is to develop the necessary trust registry server and resolver software to speak the TRQP protocol. This is underway with various software and service vendors collaborating within and across multiple industry organisations including the ToIP Foundation, the Decentralized Identity Foundation (DIF), the World Wide Web Consortium (W3C), and the OpenWallet Foundation.

Step Two: Governance

The second step is to establish a neutral nonprofit governing body for a global trust registry network—essentially the trust registry network equivalent of ICANN. This work commenced in March 2024 with the inception meeting for the "Global Acceptance Network" (aka "GAN"—now being formally branded as the Ayra Trust Network).

The first public website for the GAN went live on 17 September 2024. As it states there:

> *The Global Acceptance Network (GAN) Organization is a neutral, nonprofit entity dedicated to establishing a public utility for digital trust. Just as the Internet's Domain Name System required global, multi-stakeholder governance through ICANN, the GAN Association serves as the governing body for this new layer of digital trust infrastructure. As a governing body, the primary responsibility of the GAN Organization is to develop and maintain the GAN Governance Framework: the set of policies, rules, and technical interoperability requirements that enable GAN member ecosystems to interoperate while still independently evolving at their own pace.*

[6] This is a simplified view; in actual practice, these steps evolved iteratively, each influencing the other.

The GAN Organization—now officially established as a Swiss Association called the **Ayra Association**—will soon be publishing more information about its multi-stakeholder governance model. In this paper, we will focus on what this means for the formation of Ayra-recognized ecosystems.

Step Three: Recognized Ecosystems (Trust Clusters)

The third step is to establish the first recognized trust registries within the Ayra Trust Network. This is the work currently underway among the founding members of the Ayra Association. But again—because trust registries serve a very different function than domain name registries—this is where the ecosystem-of-ecosystems model and the DNS model diverge the most. To summarise:

- **Naming domains** are for control of pointers to Internet resources within the domain—or pointers to other naming domains. Naming domains are *hierarchical*.
- **Trust domains (**aka **digital trust ecosystems**) are for control of verification data (aka authority statements) about participants in the trust domain—or recognition of other trust domains. Trust domains are *heterarchical*.

For this reason, there are no "top level ecosystems" in the Ayra Trust Network. There are only **recognized ecosystems** that have been formed to meet the real-world trust verification needs of Ayra members. Not surprisingly, many of these ecosystems will be a fractal of the Ayra Trust Network as a whole, i.e., they too will follow the ecosystem-of-ecosystems model. Each of these collections of collaborating ecosystems are called **trust clusters**.

> See the **Trust Clusters** section of the Introduction to Ayra for an overview of the trust clusters being formed for **financial services, workforce credentials, AI agents, supply chain & sustainability, organization ID,** and **personhood credentials**.

Forming Ayra trust clusters following this fractal model has several major advantages:

1. It enables Ayra Association members who share common interests (by use case, industry sector, jurisdiction, etc.) to collaborate on their specific shared problem space.
2. It maximises scalability by spreading out the load across a decentralised global network of trust registries.
3. It avoids single points of failure.
4. It pushes trust decisions as closely as possible to the affected parties—a fundamental principle of the chaordic model pioneered by Dee Hock, the founder of Visa.

This final point means that it is up to the governing bodies of each trust cluster to decide which other ecosystems to recognize (if there is a heterarchical peer-to-peer relationship) or delegate to (if there is an hierarchical relationship).

PART THREE: The Trust Cluster Formation Process

As discussed in Step Three above, the Ayra Trust Network does not become a real functioning digital public utility until the first Ayra trust clusters are operational. However, just as the first DNS TLDs were developed more or less organically, either by private industry or public/private collaborations, the same is likely to be true of the first Ayra clusters.

In this section we discuss the role Ayra Association can play in supporting cluster formation.

> The driving force of any cluster is a set of common interests shared by multiple participants across multiple ecosystems. The evolutionary stages suggested in this section are all optional; any particular cluster could stop at any stage—or jump directly into any stage—depending on the shared trust objectives of its stakeholders.

Stage 1: Informal Cluster Discussions

In the first stage, a cluster begins as a nexus of activity among the stakeholders—any combina- tion of Ayra Association members in any membership category who have a shared interest in a set of digital trust objectives. To support this early formation stage, the Ayra Association offers two collaboration tools:

1. **Ayra Community Slack channels.** Any Ayra community member can request a new Slack channel for discussing a potential cluster. Each new cluster Slack channel will be announced via the monthly Ayra newsletter.
2. **Ayra GitHub discussion groups.** When there is enough interest, cluster discussions can advance to a GitHub discussion group where members can have multiple discussion threads with richer text. Note: Github discussions do not require technical skills; anyone who can use web mail or Google Docs can contribute to a GitHub discussion.

Stage 2: Cluster Projects

When there is enough interest, a Ayra Association member can propose to form a Cluster Project. This begins with publishing a Project charter on the Ayra wiki, sharing it via the Ayra Community mailing list, and signing up members.

A Cluster Project can specify any relevant set of deliverables such as:

1. A white paper detailing the use cases and shared objectives of the cluster.
2. Proposed credential definitions and issuer, holder, and verifier requirements.
3. Best practices for participating ecosystems, issuers, holders, and verifiers.
4. A cluster metaregistry governance framework (see Stage 3).
5. A full cluster governance framework (see Stage 4).

Stage 3: Cluster Metaregistries

The next stage in formalising a cluster is to stand up a metaregistry of member ecosystems as illustrated in Figure 7. This requires only a lightweight governance framework that simply creates a formal, discoverable association between the participating ecosystems by specifying:

1. The name and trust mark for the cluster.
2. The requirements for an ecosystem to qualify as a member.
3. The qualification and/or certification process.
4. The lifecycle requirements for metaregistry listings (e.g., application, qualification, provisioning, maintenance, suspension, revocation, reapplication).

Note that such a **metaregistry governance framework** should be kept as simple as possible, governing only what is needed for a functional metaregistry and nothing more. In some cases a metaregistry may be all the cluster needs; in other cases it may be an interim step in the evolution towards a full cluster governance framework.

Stage 4: Cluster Governance Frameworks

The final stage of maturity—for clusters who decide it is warranted—is a full-blown **cluster governance framework**. Like any ecosystem governance framework, its purpose is to establish shared credential definitions, policies, and certification requirements that apply to members of the ecosystem.

Since the members of a cluster are ecosystems, this stage implements a **layered governance model**, where the cluster governance framework defines the rules that apply at the cluster layer, and each member ecosystem governance framework defines more specialised requirements that apply within that specific ecosystem (without contravening the cluster-layer requirements).

Cluster governance frameworks provide the highest degree of interoperability and synergy between their member ecosystems. At the same time, they also require the highest degree of cohesion between the member ecosystems, and thus are likely to take longer to evolve.

To get started on an Ayra trust cluster, please connect with us at the Ayra Trust Network website: https://ayra.forum/