CARDEA

LF PUBLIC HEALTH

# The Cardea Project

## A complete, decentralized, open source system for sharing medical data in a privacy-preserving way with machine readable governance for establishing trust

Optimized for COVID-19 related proof-of-testing and vaccination

**V 1.0, June 2022**

Developed by Indicio for SITA and donated by SITA to Linux Foundation Public Health

Trevor Butterworth[1], Ken Ebert, Mike Ebert, Helen Garneau, Keela Shatzkin

*[1] for correspondence, Trevor@indicio.tech*

# Summary

Cardea is a complete, interoperable ecosystem for the creation, exchange, and verification of privacy-preserving digital health credentials using open source and open standard decentralized identity technology.

Cardea can be deployed as a layer on top of existing identity systems. This means it is easy to integrate. It is also easy to use, requiring only ubiquitous and familiar technology (Android and Apple mobile devices and apps) and simple, easy-to-follow behaviors (swiping). It is also able to function offline.

Cardea was developed and successfully trialed for proving COVID-19 test and vaccine status for travelers entering Aruba and visiting hospitality spaces across the island. It did this in a way that avoided the need for third parties to manage or store information for the purpose of authenticating the data, and facilitating the user's consent to share data at point of entry.

Cardea also successfully integrated health status data from a Health Information Exchange—New York's Bronx Rhio— in the U.S. with island travel requirements such that travelers could be approved for entry before arriving at the airport.

Based on the Hyperledger Foundation's open source Indy, Aries, and Ursa decentralized identity codebases, the code for Cardea emerged from the "Happy Traveler Card," developed by Indicio.tech with SITA, the world's leading specialist in air transport communications and information technology, and the government of Aruba.

The unique parts of the code were donated by SITA to Linux Foundation Public Health, where it is a project supported by a community of developers, including Indicio, SITA, IdRamp, and Liquid Avatar Technologies.

While COVID-19 was the impetus for Cardea, the Cardea codebase provides a flexible foundation for sharing other kinds of health data in a privacy preserving way, allowing the user to manage consent for data sharing, as delineated in the machine readable governance framework.

This paper provides some contextual background to the emergence of decentralized identity technology and the problems it solves, especially in comparison with legacy, X.509 based federated identity systems. It describes the development of Cardea, how it works, the trials in Aruba, other use cases, and the future development roadmap

# Contents

# Mission and background

The Cardea Project's mission is to provide public health authorities with a complete, decentralized, open-source system for sharing health data that preserves individual privacy and consent while enabling authoritative verification of the source and validity of their data.

Cardea was initially designed to handle COVID-19 testing and vaccination requirements for international travel in a way that mitigated the risk of fraud and avoided the need for third-parties to store data or manage verification. It was developed by Indicio.tech with SITA, the world's leading specialist in air transport communications and information technology, and the government of Aruba, as the "Happy Traveler Card" (see the section "Cardea in Practice: Aruba").

The timing was fortuitous as it aligned with the scheduled implementation in April 2021 of a data-sharing provision in the 21st Century Cures Act, a 2016 US regulation simplifying a patient's right to their health data, originally provisioned in a 2000 Privacy Rule of the Health Insurance Portability and Accountability Act (HIPAA) of 1996.

Under the Cures Act, patients must have direct digital access to eight categories of clinical notes in an electronic health record, including lab test results. This regulation encourages the development of solutions like Cardea by requiring healthcare organizations and data holders to unlock access to a patient's data.

The result of this access is that the patient gains control of their data and Cardea allows them to manage it in a secure and verifiable way.

After successfully managing COVID-19-related travel requirements, and in line with Indicio and SITA's commitment to open-source technology, the Cardea codebase was donated by SITA to Linux Foundation Public Health (LFPH) in July 2021 so that public health authorities everywhere could avail themselves of this technology. Indicio's CTO Ken Ebert and Keela Shatzkin of Shatzkin Systems lead the Cardea Community Group at LFPH to continue development of the codebase to incorporate additional features and requirements.

While Cardea is designed to provide individuals with the tools needed to safely share COVID-19 test results and vaccination details via Android and Apple mobile devices, the ongoing development of Cardea and the growth of the Community Group reflects Cardea's capacity to adapt and function as a platform for sharing and verifying multiple kinds of health data, through privacy-by-design and security-by-design principles.

As such, managing COVID-19 testing and vaccination data is only the first of many use cases that range from patient data sharing to prescription management (see the section Use Cases Beyond COVID-19).

## Why open source?

Open source technologies drive a virtuous cycle of innovation.  As implementations increase, the developer community grows, which in turn feeds more implementation and innovation. Indicio, SITA, and other participants in the Cardea Community Group see open source as essential for scaling Cardea across multiple partners, sectors, and countries, and as a key to sustainable evolution.

Similarly, open source is essential to global inclusion and equity. The technology behind Cardea can deliver significant public benefit through the creation of Trusted Digital Ecosystems that enable consent-based data sharing, privacy protection, and a level of security impossible in centralized and federated identity systems. These benefits should be available to everyone, and donating the code to Linux Foundation Public Health means these benefits are available to everyone.

## Why decentralized identity?

Decentralized identity reinvents the way digital identity is created and used. In doing so, it solves a fundamental problem in the architecture of the internet: There is no way to create, own, and assert a unique digital identity if you are a person.

Before the advent of the web, individual identity on the internet wasn't a problem. The internet—then known as ARPANET—was a way to communicate between computers (with the objective of networking distributed computational resources to execute memory-intensive tasks); computers could be found through their Internet Protocol (IP) addresses, each designated by a unique series of numbers.

With the creation of the World Wide Web, websites—ways of structuring information that could be searched and accessed through a web browser—also acquired identities in the form of web addresses (Uniform Resource Locators—URLs), each unique and based on the Domain Name System governed by a global organization, ICANN.

Human identity emerged online through the use of email addresses as identifiers. In broad terms, these identities were not asserted but leased from a company or institution acting as an email provider. They required creating a profile, which included a name (which, importantly, did not need to be your legal name) and a password to login and access messages to your address.

Email is a form of centralized identity. It is not fully portable, as you cannot always take your email address to another provider and service; it is not permanent, as you can have your account revoked by the provider and lose all the information you have accumulated through the use of this email identity; and your profile information and all the data associated with it have to be stored in a database by the provider in order to make the verification of your identity possible.

This centralized model of a user profile, login, password, and other personally identifying information (PII) has become the basis for facilitating human identity online, for accessing services, for ecommerce, for work, and for using social media.

## The problem with centralized authentication

Centralized identity architecture means that a person must replicate their user profile—along with the requisite PII—across multiple platforms to access multiple services.

As each of these digital identities functions as a key to enter a database, a single compromised identity can, when it is used to enter a company network, allow access to the entire network. This means that each digital identity in a database is a potential single point-of-failure for the entire network.

At the same time, the combination of increased digital banking services and e-commerce have increased the risk of stolen PII being used to commit fraud, either on its own terms or in combination with other, publicly accessible, data.

For commercial services, where access to the product or service requires creating a user profile using a leased email identity, the aggregation of user PII in a siloed database also makes that database an attractive target for cybercrime.

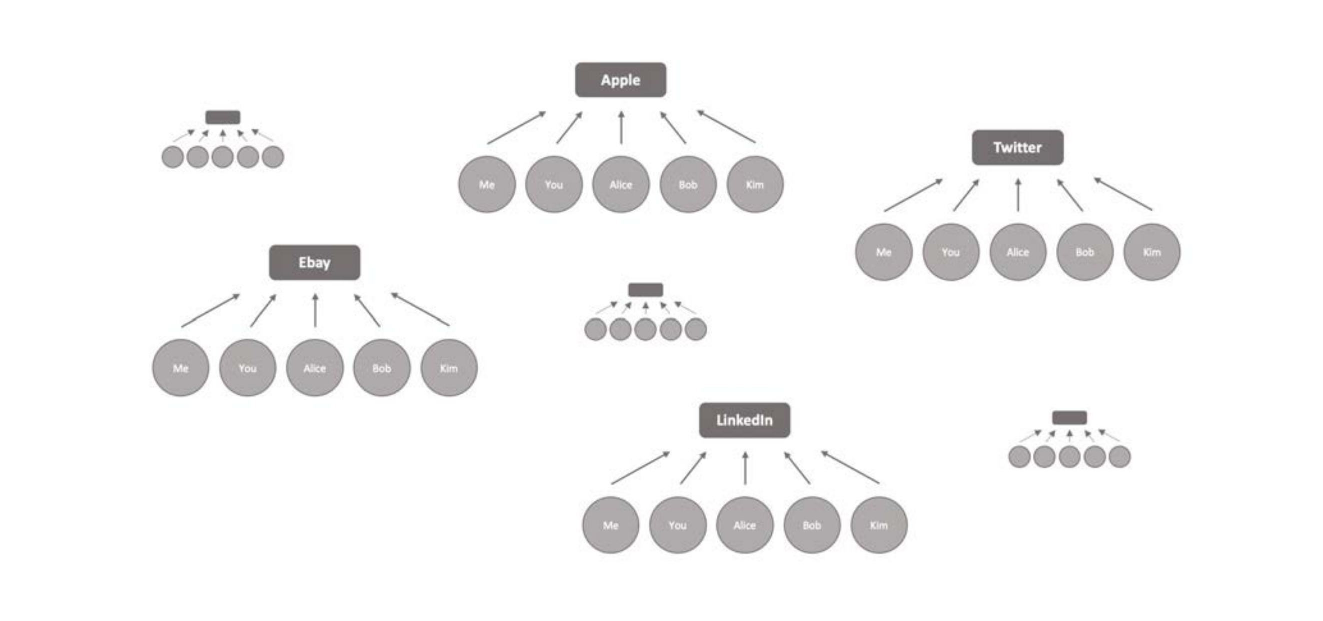

## Federated identity: A new model with new problems

Federated identity adds limited portability to a centralized identity: An authenticated login to a user profile on one platform enables access to other associated platforms and services. To some degree this simplifies the complexity of managing dozens of passwords and logins for web services; but it also means that multiple services are dependent on the integrity of a single digital identity. If the platform is compromised, access to all those services becomes impossible.

At the same time, many digital identity providers extract a form of rent from each of the identities they provide: Users must agree to allow their behavior online to be tracked in exchange for their digital identity. They may also have to agree in advance to this metadata being sold to other parties of the identity provider's choosing at any time. The more user metadata is aggregated, the more valuable it is as a monetizable asset to an identity provider; the more valuable it is, the greater the risk of it being stolen.

These synergies have driven an explosion in identity theft and cybercrime: 92 percent of all malware is delivered by email; the estimated global cost of cybercrime will be $10.5 trillion by 2025. Healthcare is increasingly a target: 93 percent of healthcare organizations have reported a data breach, there were 154 ransomware attacks on healthcare in 2021, and the the U.S. Department of Health and Human Services' Health Sector Cybersecurity Coordination Center has issued a threat brief on the specific vulnerability of EHRs, as they contain the PII to enable identity theft and tax and insurance fraud.

Solutions for mitigating these threats include hardware security tokens, two-factor authentication (2FA), SMS 2FA and software one-time passwords. In 2013, Apple took biometric authentication to the masses with Touch ID. Yet none of these advances, many of which are cumbersome for users, have replaced dependence on the password. Instead, they are often used alongside a password either protecting or adding a shared secret. Worse, even where enhanced security features are provided, consumers will choose less secure alternatives if those alternatives are frictionless.

The combination of security failures and privacy concerns has resulted in stringent and complex data privacy law, notably the European Union's General Data Protection Regulation (GDPR); it has driven increased public concern about predatory behavior by technology companies, captured in the phrase "surveillance capitalism;" and it has also driven a fundamental rethink in how digital assets should be secured, leading to the rise of a new, dominant paradigm in cybersecurity, "Zero Trust."

Zero Trust emerged with the realization—based on the analysis of thousands of data breaches—that everything behind a network perimeter was at risk from a single breach of that perimeter. This made every user, device, and access point a security threat if their identity and access credentials were stolen. Therefore, network security should be managed under the assumption that the network has been breached; user verification must be continuous for interaction. In many ways, the concept of Zero Trust is a fitting summary for the state of online interaction;

it has reached an impasse, unable to provide the necessary authentication, privacy, and security needed by its participants from within its own architecture.

The challenge, especially in federated identity systems, is that the inherent Zero Trust principle of “I don’t trust you and you don’t trust me” is uni-directional. Organizations must still trust the federated identity service provider, and the individual must trust both the service provider and the system they wish to connect to. In solving this problem, decentralized identity is naturally the “other side of the coin” to Zero Trust architecture.

## A better architecture for creating trust in data and in digital relationships

Decentralization is a foundational concept to creating a trusted digital ecosystem such as Cardea. In simple terms, it means replacing the centralized control of identity for managing how we authenticate data. This, in turn, means no longer having to rely on user profiles, passwords, logins, and PII to establish connections and authenticate people and information; it means there is no need for a third party to store that personal information to facilitate verification.

Instead, any governing authority—a government, an institution, an association, a company, or a store—can issue a digital credential whose authenticity and ownership can be verified without having to check against the contents of that credential or check in with the original source, based on an agreed-upon trust framework and governance rules for sharing and verifying data.

*It is possible to directly establish the uniqueness for a collection of data without having to rely on the data or a broker of the data to provide the proof of its authenticity*

Whether the issuing entity is trusted for a transaction is a separate matter: A verifiable store card may not cut it for a mortgage lender as a proof of identification compared to a verifiable bank card (with all the real world assurance proofing that entails); but each card is uniquely verifiable on its own terms.

This has significant implications for managing privacy and improving security. It is possible to directly establish the uniqueness for a collection of data without having to rely on the data or a broker of the data to provide the proof of its authenticity. This is also the basis for the European Union’s e-identity wallet model as part of the eIDaS2 legislation.

This also means that if the data inside a credential is shared, such as a traveler presenting a particular COVID-19 test taken on a particular date to a border agency, that data is verifiable by the "material" fact of the credential surrounding it being independently verifiable. We can know that it was issued to the traveler by an entity that can be verified against an approved trust registry of verifiable issuers.

*This technology is extraordinarily flexible and powerful, which is why it is often difficult to easily explain. It provides new pathways for managing authentication, privacy, and security all at the same time—and these can be combined in different ways to solve different problems*

This is because the parameters governing a verifiable credential—the information describing its source, the kind of information it contains, and who it has been issued by—are anchored to a blockchain-based distributed ledger. The ledger-based information cannot be changed with modifying the blockchain; as blockchains are immutable, this would break the chain and be detected. As a result of cryptographic signing, the credential-based information cannot be changed without detection. Finally, and importantly, no personally identifiable information (PII) is written to the ledger.

As the ledger is distributed across a network, we also avoid the risk of a single database becoming a point of failure or a federated identity system going offline: There are multiple copies of the information governing the verification of a verifiable credential across the globe (for example, the Indicio Network, a Hyperledger Indy-based network for digital identity has multiple "nodes"—copies of the ledger—on five continents).

These capacities for privacy are enhanced by the way a person can hold their information and how they can share it from a verifiable credential.

Depending on the format of the credential, a person can select the information they wish to disclose instead of—as is typical with analogue ID—showing all their personal information in one go.

It is also possible to use what are called "zero knowledge" or "predicate" proofs, such as proving legal age without having to disclose age, based on an accepted and verifiable cryptographic "proof" in the credential.

This technology is extraordinarily flexible and powerful, which is why it is often difficult to easily explain. It provides new pathways for managing authentication, privacy, and security all at the same time—and these can be combined in different ways to solve different problems.

Additionally, one of the communications protocols that enables these interactions (W3C DIDComm) further enlarges the scope for information sharing in ways that extend privacy and security and add semantically rich interaction. Market implementations are only beginning to explore this new communication dimension.

# Trusted Digital Ecosystems

The terms “decentralized identity,” “distributed identity,” “self-sovereign identity,” and “reusable identity” have been used to summarize all these capacities but they have, for many not immersed in the technology, resulted in a much narrower perspective; namely, that the sum of the technology is just a new kind of driver’s license.

At the same time, some in the decentralized identity community take a maximalist position, interpreting self-sovereign identity as the capacity for individuals to have complete control over their digital identities — hence, “self sovereign.”

Both perspectives miss the way that “identity” in a decentralized, verifiable credential system can mean the authentication of *any* digital information; they miss the way the technology is flexible so as to be governable by existing sovereign entities and in hierarchical ways; and they miss the urgency of decentralizing identities for machines, the internet of things, and even non-digital objects in the emerging spatial web.

The web and digital life are evolving rapidly. The digital twins, smart cities, and robots of the near future will all require identities that can be trusted. We know—or should know—from science fiction why these identities should not be centralized.

Because the terms “decentralized identity” and “self-sovereign identity” limit people’s perspectives on the technology’s capabilities, we use the category of “Trusted Digital Ecosystem” to express the interaction of a network, software agents, machine-readable governance, digital wallets, and verifiable credentials as a system.

Cardea is a Trusted Digital Ecosystem because it provides a complete, end-to-end solution to sharing and verifying health data in a privacy preserving way—and by doing that, all interactions within this ecosystem are able to be trusted.

*The web and digital life are evolving rapidly. The digital twins, smart cities, and robots of the near future will all require identities that can be trusted*

## The roles within a Trusted Digital Ecosystem: Issuer, Holder, Verifier

There are three roles in a Trusted Digital Ecosystem:

**Issuers**—those that issue verifiable credentials, such as a healthcare provider

**Holders**—those who are usually the subjects of the verifiable credentials—patients, employees, or travelers—and hold them on a mobile device (the holder may also be a legal guardian of the subject, such as a caregiver or parent)

**Verifiers**—those who need to verify the information contained within these credentials, such as a hospital, a health or relief agency, or an event space or a hospitality venue.

Each of these roles has its own designated software, i.e., for issuing a credential, for holding and presenting a credential, and for verifying a credential.

We call these software "agents" because they manage the information flow between parties and each has a fiduciary responsibility to the party it represents.

There are also mobile and cloud-based agents. Mobile agents enable connections with mobile devices (which lack fixed IP addresses), while cloud agents enable more advanced automation and workflows and can act as backups in the event someone loses their mobile device. (Note, these are not currently part of Cardea's reference implementation but are on its development roadmap.)

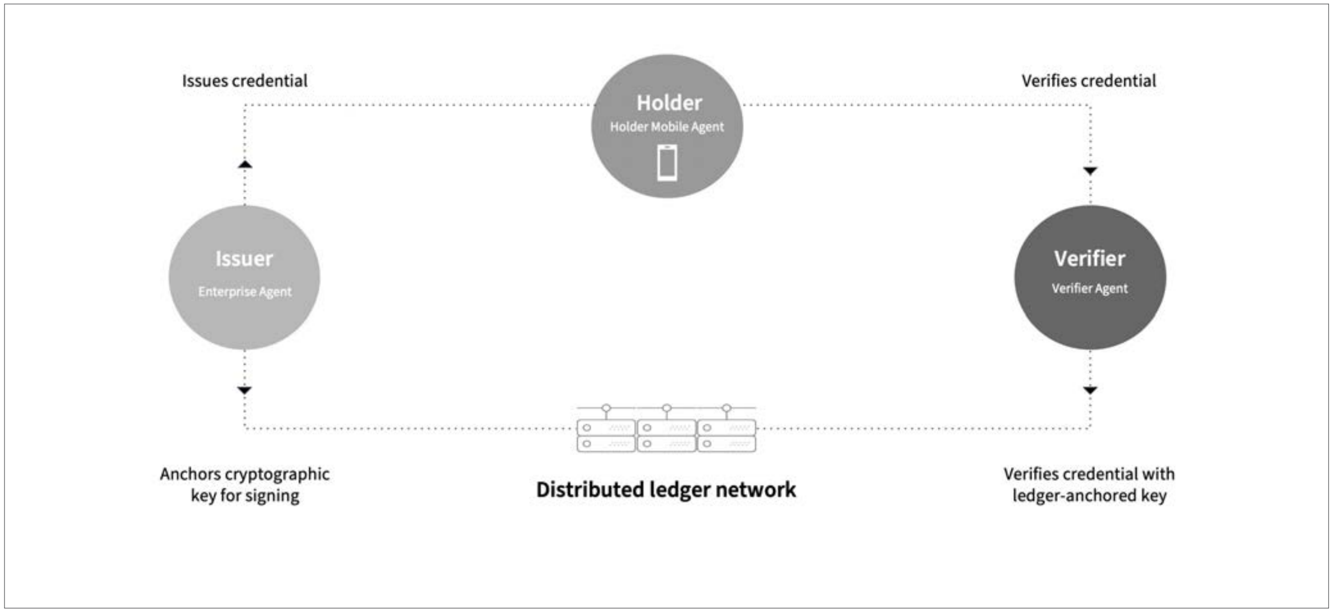
Issues credential
Holder
Holder Mobile Agent
Verifies credential
Issuer
Enterprise Agent
Verifier
Verifier Agent
Anchors cryptographic key for signing
Distributed ledger network
Verifies credential with ledger-anchored key

## Machine readable governance

Agents also reference machine readable governance files to encode the decisions normally made by humans in a machine readable format.

Machine readable governance files are published by the governance authorities for an ecosystem and distributed to all the agents in that ecosystem. These files are an efficient and error-free way to choreograph the rules for interaction in a jurisdiction: For example, if a traveler has a COVID positive test that's more than two weeks old, they would be considered recovered and allowed in.

Similarly, if they have a negative COVID test within a designated lookback period, let's say 24 hours, that can also be evaluated and determined to be acceptable for entry. Leaving these assessments to a human means possible error in calculating the date/time differentials.

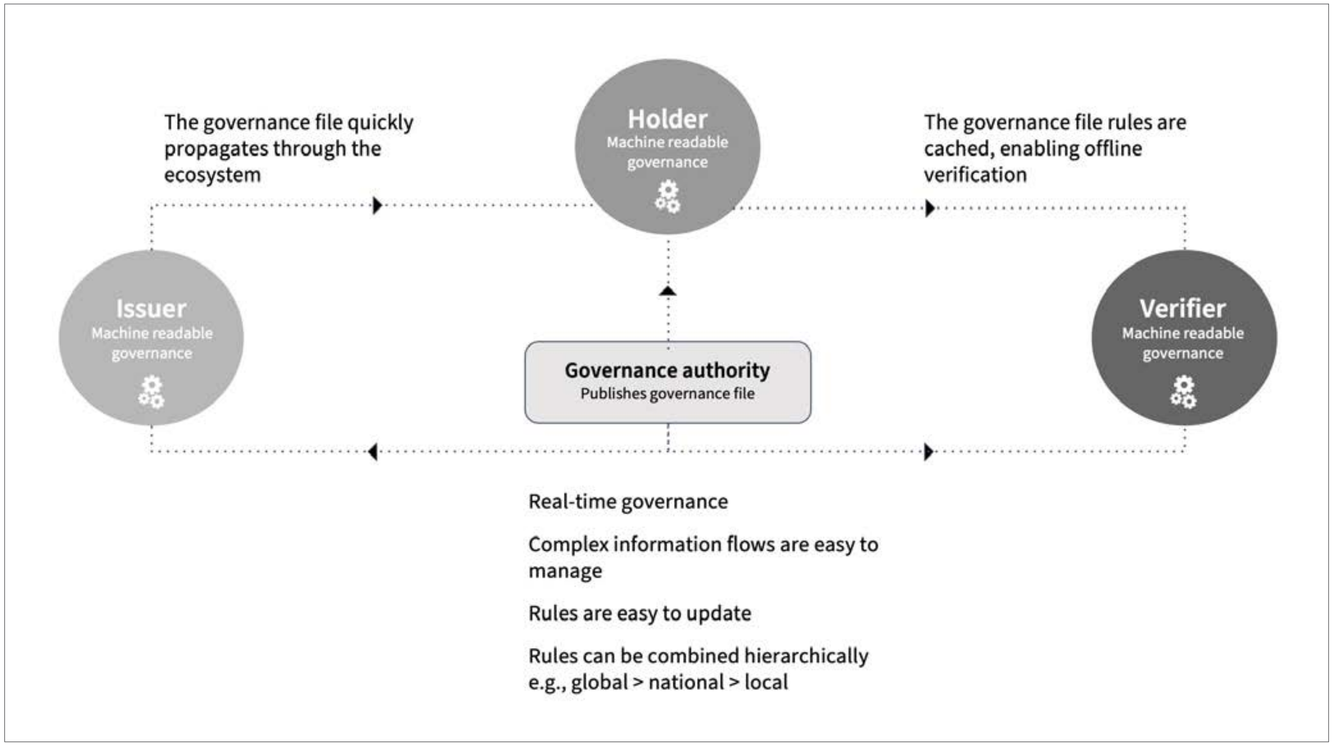
The governance file quickly propagates through the ecosystem
Holder
Machine readable governance
The governance file rules are cached, enabling offline verification
Issuer
Machine readable governance
Governance authority
Publishes governance file
Verifier
Machine readable governance
Real-time governance
Complex information flows are easy to manage
Rules are easy to update
Rules can be combined hierarchically e.g., global > national > local

The magic of machine readable governance is that these rules can easily be updated across the system as the science or regulations change, without risking that someone missed that memo.

As these files are held at the agent layer, they are cached and therefore can function offline. Machine-readable governance makes a governance framework portable; instead of a centralized trust registry to verify issuers in real time, a governance authority publishes a governance file that propagates to all the agents in a Trusted Digital Ecosystem.

With Cardea, machine readable governance enables the appropriate governance authority—such as a government or health authority—to directly implement and easily update the applicable rules for how individual data is used.

Machine readable governance files can also manage interactions within and across jurisdictions and health care systems, allowing each governance authority to enact the rules they decide on as important.

## Interaction

DIDS—agents enable interaction within a Trusted Digital Ecosystem by creating Decentralized Identifiers, or DIDs.

If a device has an IP address and a website, a URL or web address, the digital identity of someone or something begins with a unique DID. A DID is a URI (a uniform resource identifier), meaning it identifies some resource.

DIDs are relatively new, and their specification is governed by the W3C. What makes DIDs different from IP addresses and URLs is that they are not leased from a third party, managed by a third party, or are reusable; instead, anyone with the appropriate agent software can create a DID and there is no limit on their number.

When a software agent creates a DID for someone or something, it creates a related DID document. This DID document contains information about how to interact with the someone or something in control of the DID.

For example, if a hospital issued a health credential, the DID document would contain the mechanism by which the hospital proves it created the DID for the credential and the information to connect with the hospital to receive the credential.

This is able to happen through the use of public cryptographic keys. Parties in Cardea connect with each other by using their private keys to encrypt the information they send. Paired with a public key, they enable peer-to-peer encrypted communication.

Issuers use public DIDs so that the necessary cryptographic and structural  information for verifying a credential can be found on the ledger. No personally identifiable information is written to the ledger for verifying a credential; personal data is sent directly to the Holder in a verifiable credential.

Mobile Agents create  unique, private  DIDs for every communication channel. This makes DIDs from a single entity non-correlatable: There's a new and unique DID for every interaction.

DIDs have their own communications protocol, unsurprisingly called DIDComm. This can be accessed over any transport and allows for rich, semantic interaction and offline verification directly between devices.

DIDs are the building blocks for verifiable identity credentials, such as  passports, diplomas, or a COVID-19-test result.

### Schemas and Credential Definitions

A schema is a specification for verifiable credentials that defines and describes the data they contain.

A credential definition describes the form of the credential, associates the schema with a particular issuer, contains the cryptographic material to support selective disclosure, and establishes an optional revocation mechanism (for listing the revoked credential in a registry).

No individual content or PII is written to the ledger.

The Cardea ecosystem supports the concurrent use of multiple schemas and credential definitions.

### Revocation

Credentials must be revocable, whether due to errors in issuance or breach of user terms. A Revocation Registry provides a privacy-preserving way to do this.

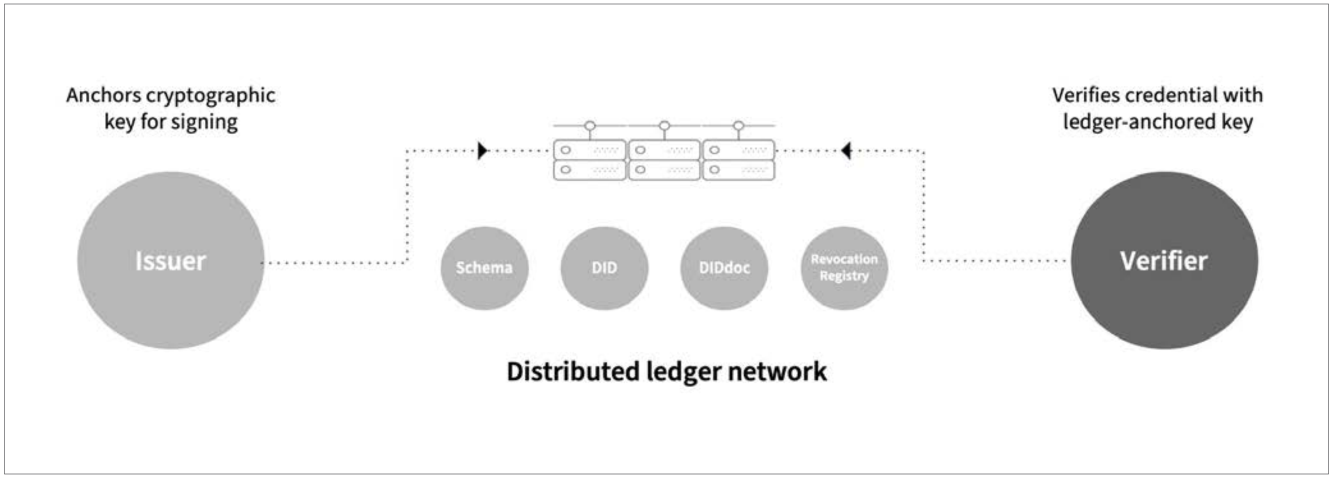
Anchors cryptographic key for signing
Verifies credential with ledger-anchored key
Issuer
Schema
DID
DIDdoc
Revocation Registry
Verifier
Distributed ledger network

## Authentication in a Trusted Digital Ecosystem

The user experience of these mechanisms is maximally frictionless. Once past the identity assurance systems used by a lab or a hospital to create a patient record, authentication in a Trusted Digital Ecosystem requires no more than scanning, tapping or  swiping to consent, then sharing, and verifying.

1. The Issuer writes a public DID to a distributed ledger so that a person's software agent can contact the issuer to receive a credential.

2. The Issuer adds a credential schema and definition to the ledger. This describes the form of the credential, associates the schema with a particular issuer, and establishes an expiry date or revocation mechanism (for listing the revoked credential in a registry if terms of use are violated). No individual content or PII is written to the ledger.

3. The Issuer emails or displays an invitation to a person so they can connect and receive the credential. If accepted, the issuer will generate a unique public DID to directly communicate with the person via their application. If accepted, the holder will generate a unique  DID to communicate with an issuer.

4. Public and private keys are generated within this connection for verifying identity and for the encryption and decryption of communications.

5. The Issuer creates a credential for the person (following the Issuer's normal assurance or "know-your-customer" (KYC) process for identity proofing).

6. The person accepts the credential, checks that the information is correct (there is a step to correct any errors), and then stores the credential in a software agent (if they don't have a software agent, the Issuer will prompt them to download one). Software agents are contained in standalone mobile apps (or integrated with existing apps), or they can be hosted in the cloud.

7. There are two ways to use the credential depending on the privacy required. Either a person can present a proof of the credential so that their identity, or a fact related to their identity, can be verified without the need to share any specific information, or the person can respond to requests for specific information. In this second case, they also have the option of sharing some specific information in a privacy-preserving way, using selective disclosure and predicate proofs powered by zero-knowledge credential technology.

8. The Verifier is able to check the authenticity of the credential and its ownership by looking up the Issuer's public key and DID on the network and validating the credential. This means the issuing authority doesn't have to be consulted to prove the authenticity of the credential. This also removes the need for a third party to facilitate verification by using PII stored for cross checking or the need for a direct integration between the verifying organization and the issuing organization.

A Trusted Digital Ecosystem means that *each* interaction in an information flow is a uniquely encrypted, non-correlatable, non-reusable, peer-to-peer interaction. When combined with the privacy-preserving features of a credential (selective disclosure and zero-knowledge proofs), a Trusted Digital Ecosystem provides a new, multi-dimensional level of security for data sharing and verification. When combined with machine readable governance, dynamic rulesets, based on geography, regulations, and data type, can be applied to further facilitate secure and private data sharing.

## Issuing, presenting, and verifying in Cardea

A Cardea ecosystem contains the following components:

- A Health Enterprise Agent and Schema(s)
- A Mobile and an Enterprise Verifier Agent
- A Holder Mobile Agent
- A Hyperledger Indy Network

- A Government Enterprise Agent and Schema (if following the information flow in Aruba, where the government issued a derivative credential from the test credential).

An **Issuer**, such as a laboratory, a healthcare provider, or a state health agency, uses a **Health Enterprise Agent** to write a **public DID** to a **Hyperledger Indy Network** announcing themselves as an Issuer.

It then issues a cryptographically-signed digital credential using a defined **Schema**, a template that allows all relevant parties to know what to expect in terms of the structure and content of information inside a credential. In Cardea, current Schemas are neutral and support any kind of lab or vaccine content or health verification scenario.

**Cardea information flow 1**
**Health provider creates verifiable credential**

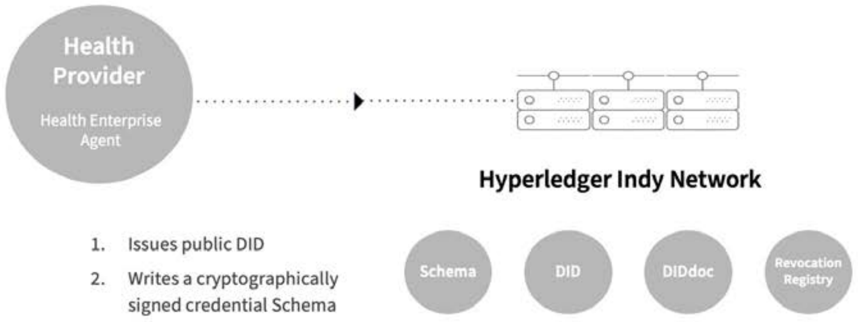

1. Issues public DID
2. Writes a cryptographically signed credential Schema

## Accepting a credential

This credential is accepted by a **Holder**, such as a patient or traveler, who holds it in a **Digital Wallet** on a mobile device through a **Holder Mobile Agent**. The credential contains the health information specific to its purpose (details of a test, vaccination, or exemption), and the holder is free to use the credential and its content as they see fit.

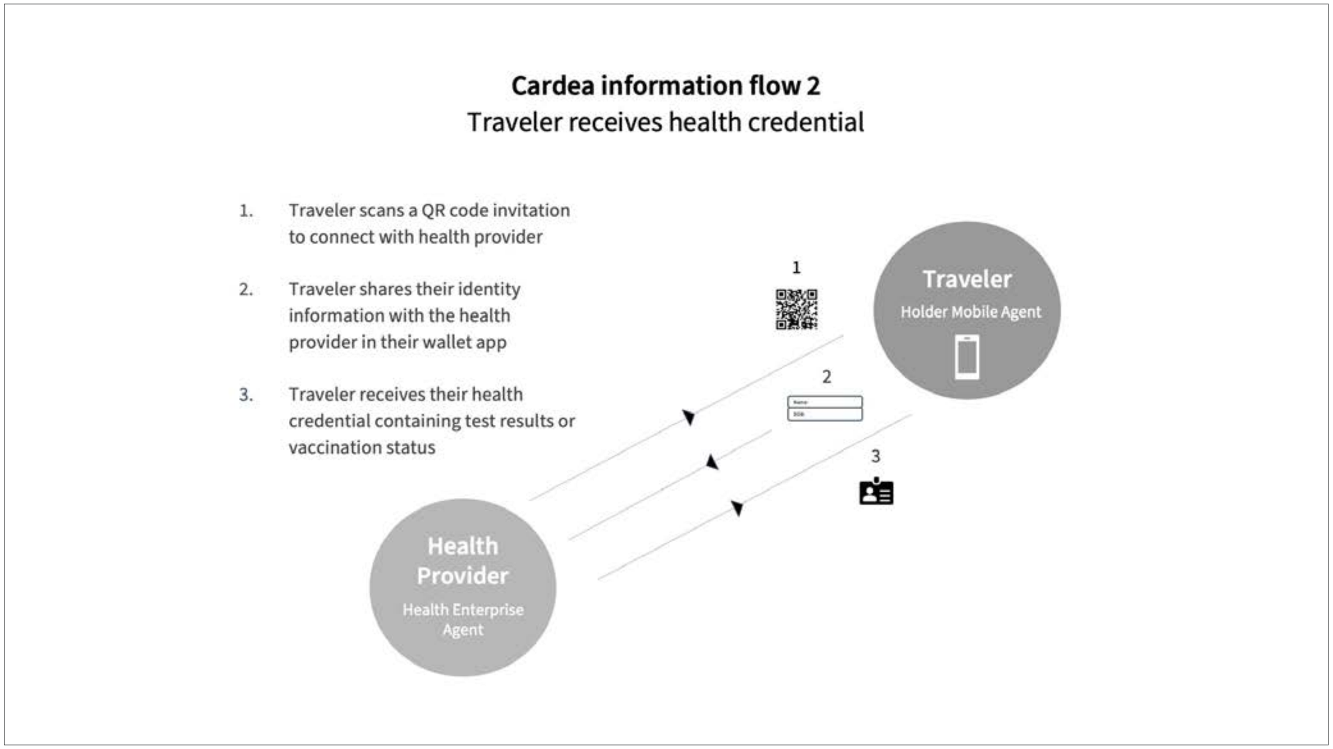

**Cardea information flow 2**

**Traveler receives health credential**

1. Traveler scans a QR code invitation to connect with health provider
2. Traveler shares their identity information with the health provider in their wallet app
3. Traveler receives their health credential containing test results or vaccination status

## Verifying a credential

Typically, the Holder will present the credential to a **Verifier**, a verifying organization or entity that requires proof of the contents of the credential, such as proof of a test or vaccine result.

As there are privacy features built into the credential, specific information can be shared in a controlled way. This means that the Verifier sees only what they need to know for the given use of the credential. In Cardea, this is achieved through using **selective disclosure** (the holder selects and shares only the specific information needed) or **predicate proofs** (the holder is able to generate a proof without disclosing the actual information).

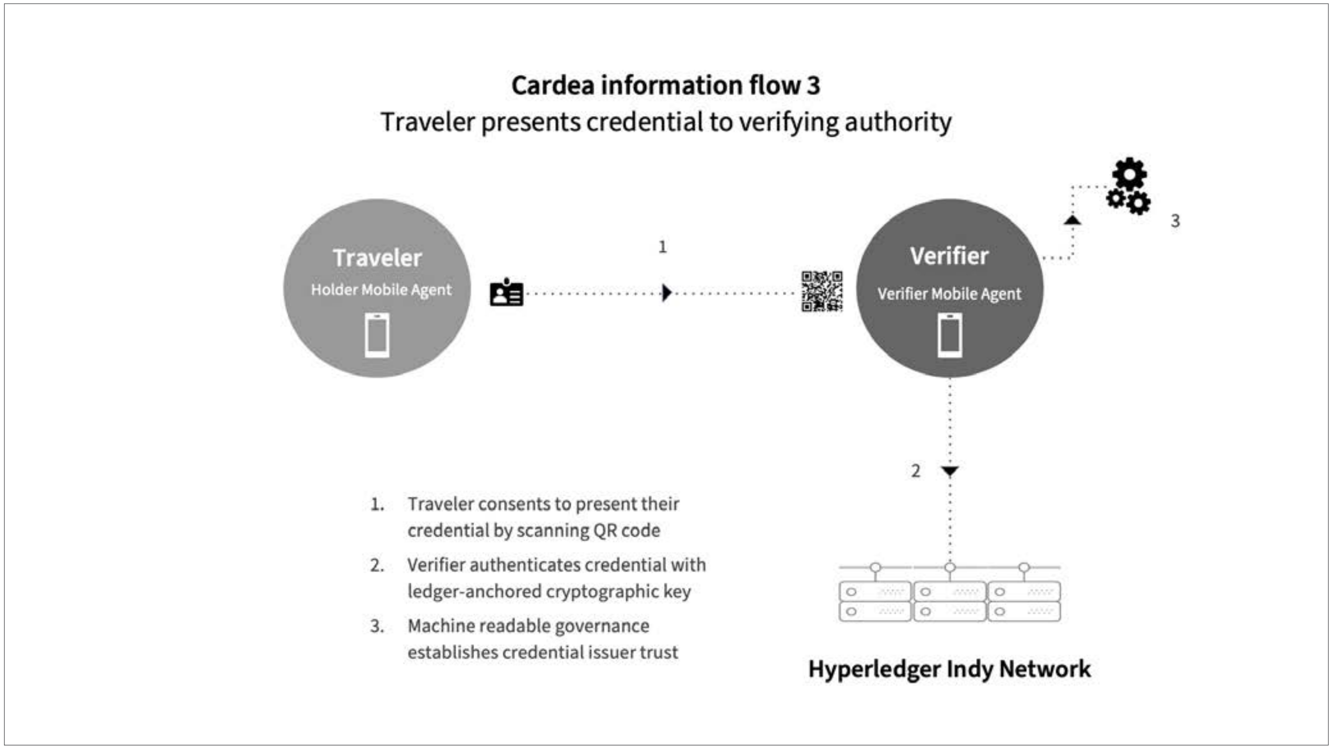

**Cardea information flow 3**
Traveler presents credential to verifying authority

1. Traveler consents to present their credential by scanning QR code
2. Verifier authenticates credential with ledger-anchored cryptographic key
3. Machine readable governance establishes credential issuer trust

## Implementation in Aruba

In the Cardea implementation in Aruba, verification was a two-step process. Upon arrival on the Island, travelers took a COVID-19 test and were issued a credential for this test through the testing lab's **Health Enterprise Agent**. The traveler used this credential to receive a test results credential.

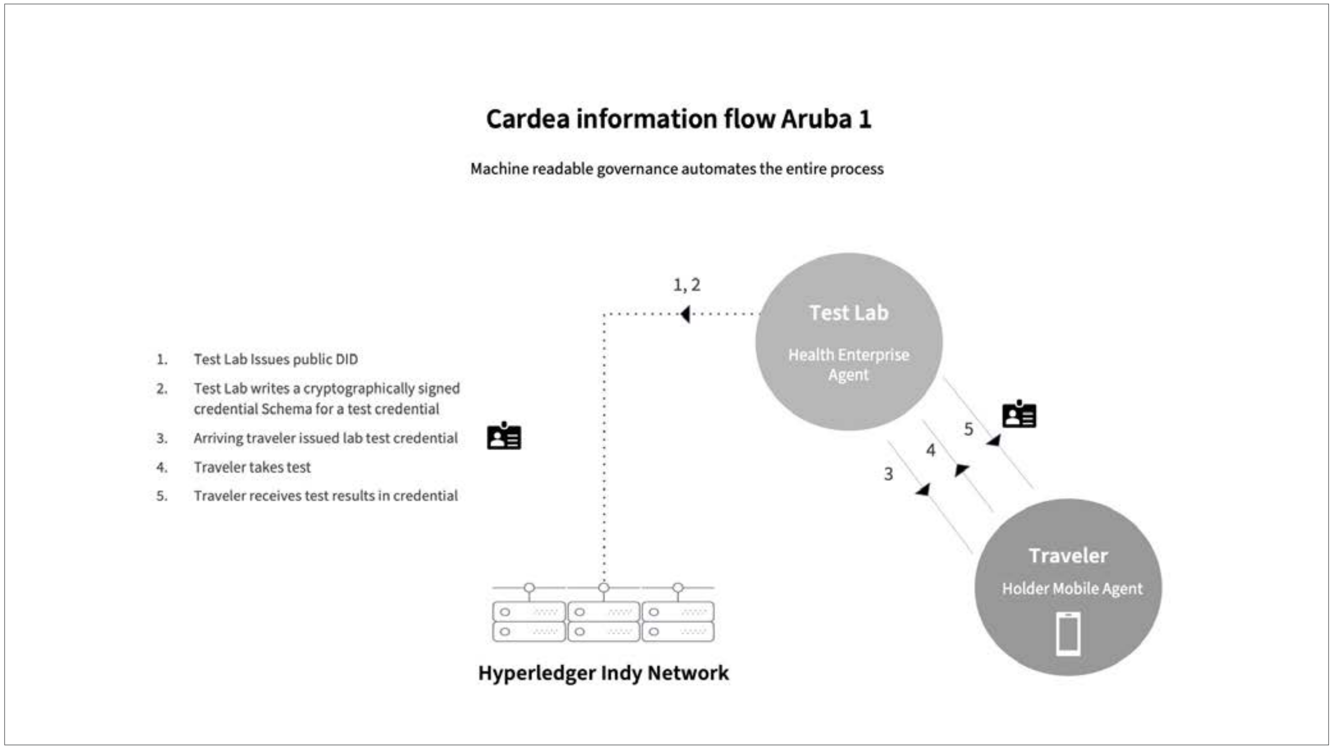

### Cardea information flow Aruba 1

Machine readable governance automates the entire process

1. Test Lab Issues public DID
2. Test Lab writes a cryptographically signed credential Schema for a test credential
3. Arriving traveler issued lab test credential
4. Traveler takes test
5. Traveler receives test results in credential

## Implementation in Aruba

The travelers shared their test results credential with the Aruban government, which authenticated its validity by verifying the Issuer’s DID. If the test was negative, the Aruban government then issued a derivative credential through a **Government Enterprise Agent**.

This derivative credential proved the traveler had tested negative but contained no health data (thereby

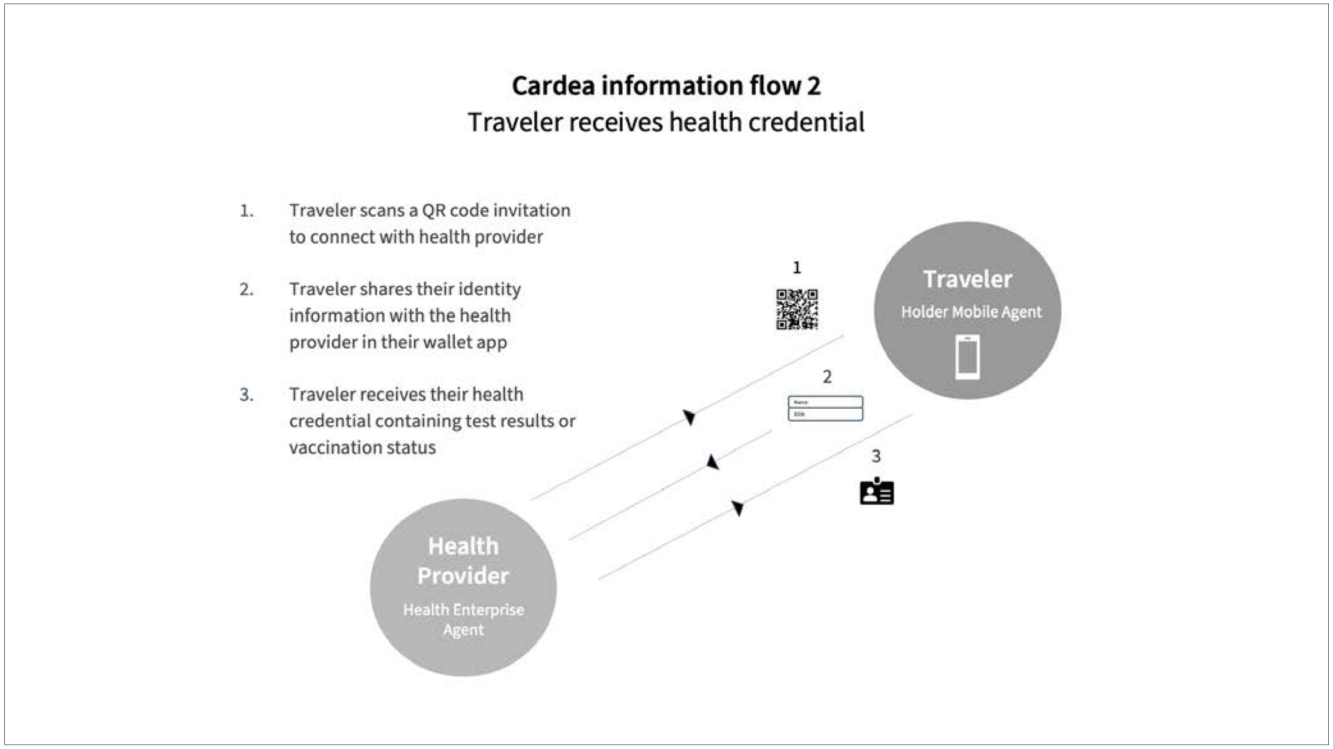

**Cardea information flow 2**

**Traveler receives health credential**

1. Traveler scans a QR code invitation to connect with health provider
2. Traveler shares their identity information with the health provider in their wallet app
3. Traveler receives their health credential containing test results or vaccination status

## Implementation in Aruba

adding an extra layer of privacy).

When the traveler presented this derivative credential at hospitality spaces around the island, it was scanned using a **Verifier Mobile Agent**, which verified the validity of the credential through the **Issuer's DID**.

During this process, communication between the Issuer and Holder and the Holder and Verifier is direct and conducted through **DIDComm**, a cryptographically secure channel that sits on top of any transport connecting the participants. DIDComm enables direct, peer-to-peer interaction, which makes interaction between participants resistant to surveillance.

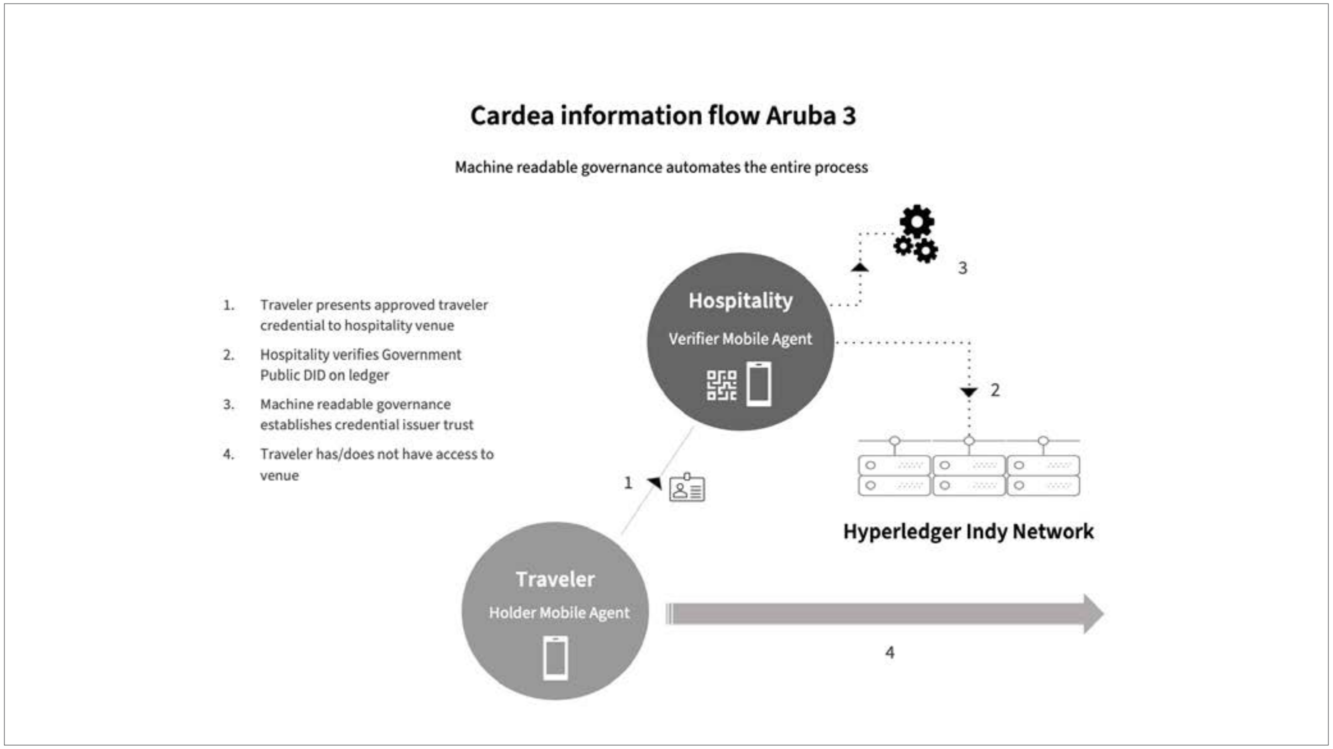

**Cardea information flow Aruba 3**

Machine readable governance automates the entire process

1. Traveler presents approved traveler credential to hospitality venue
2. Hospitality verifies Government Public DID on ledger
3. Machine readable governance establishes credential issuer trust
4. Traveler has/does not have access to venue

## Governing the information flow

Machine readable governance is a way to encode decisions normally made by people and apply rules in a format that can be read by software. For Cardea, the relevant governance authority publishes a central file which is referenced by all agents in the ecosystem.

Machine readable governance supports automated decision trees that establish (1) who has what role and (2) what happens when something passes or fails in the information flow. As these rules are predefined, governance works offline, which is critical for real-world applications.

When changes are made to the governance file, the agents can retrieve the updated file and adjust their behavior without having to redeploy server code or release mobile app updates. This means that issuers can be added or dropped, and rules can be changed rapidly as new information becomes available.

Machine readable governance is a flexible way to manage complex information flows across different ecosystems where decisions need to be (1) made locally by the presiding jurisdiction or (2) arranged hierarchically within jurisdictions.

*With Cardea, machine readable governance enables the appropriate governance authority—such as a government or health authority—to directly implement and easily update the applicable rules for how individual data is used*

## Privacy, fraud, and paper

Cardea transfers privacy control to the holder of the data as the Holder decides with whom they share their data. This addresses and simplifies the complex rules governing the sharing of personal health data.

The challenge of managing travel during Covid was not just a digital problem, it was also a paper problem. Proof of Covid testing and vaccination, whether for travel or not, relied on paper-based credentials.

As Bloomberg Businessweek put it, "All it takes to make your own vaccination card is a few minutes of online searching, a printer and some card-stock paper." If this was too much effort, the Guardian reported that "A hidden pandemic market advertising fake vaccine and test certificates for as little as £25 has grown exponentially, with more than 1,200 vendors in the UK and worldwide, researchers have found." While paper-based credentials were all too easy to produce, they were exceedingly difficult to verify, particularly when manual verification only added to airport lines.

Verifiable credentials eliminate paper forgery. They are tamper resistant due to the combination of blockchain immutability and cryptographic signing. The former ensures the credential cannot be copied or altered; the latter means that the information in the digital credential cannot be altered without detection.

The authenticity of a digital credential from a trusted issuer can, therefore, be authoritatively verified — and by using automated scanning equipment, it has none of the friction of manual verification.

The proliferation of Covid paper credentials in the face of such flaws signals the inadequacy of legacy digital identity when it comes to personal data. Many governments did not want to push people toward using commercial, third-party vendors where their health data would be collected and potentially correlated.

## How to use the Cardea codebase

Cardea is currently fully functional and deployable. The roadmap is to further enhance the codebase with additional features and functionality. Please visit Cardea.app to learn how you can connect with the Community group and get started.

This white paper is necessarily simplified. While it is not essential to understand how the technical code behind Cardea is configured and implemented to use Cardea or any Trusted Digital Ecosystem, we encourage those interested in learning more to explore the work being done at the World Wide Web Consortium (W3C), Decentralized Identity Foundation (DIF), and Hyperledger Foundation to develop and standardize the technology. Indicio also has workshops on decentralized identity for all technical levels.

## Cardea in practice: Aruba

The impetus to build out Cardea came from, and was supported by, SITA, and the code developed by Indicio, in tandem with SITA, was donated to Linux Foundation Public Health, where it became the starting point for the Cardea project.

The global pandemic showed why the ability to share and verify health data was critical to sectors beyond health care; it also showed that the ability to do this in a privacy-preserving way and without direct integrations was legally and technically imperative.

SITA also arranged for the first real-world demonstration with the assistance of the government of Aruba, including the Public Health Department, and Bronx RHIO, a New York Health Information Exchange (HIE).

The first trial was in April 2021 and tested the issuance of a health credential and derivative “Trusted Traveler” credential, which followed the information flow described above.

Indicio and SITA were able to demonstrate that (1) a credential issued by a medical provider was verifiable by the government and (2) the government’s derivative credential could be verified by hospitality venues across the island.

The second trial was held in December 2021 and focused on implementing machine readable governance. The systems assessed who could enter the island based on rules for which COVID-19 tests and vaccines were valid and were taken within specified time ranges.

Notably, instead of being tested upon arrival in Aruba, travelers were able to test in the US and Canada before they arrived. They were, therefore, able to be verified for entry to Aruba before they even got to the airport.

On May 11, 2022, SITA won the Verifiable Credentials and Decentralized Identity Award at the European Identity and Cloud Conference, hosted by KuppingerCole, for its implementation of decentralized identity in Aruba. SITA was chosen based on their innovative and unique application of verifiable credentials, the ability of the product to scale, its complexity, and the completed status.

Verifiable Credentials & Decentralized Identity
eic
european identity and cloud conference 2022
NHS
SITA
kuppingercole
celebrating 15 ye

Adrien Sanglier accepts the Verifiable Credentials & Decentralized Identity award on behalf of SITA at the KuppingerCole Analysts AG European Identity and Cloud Conference, Berlin, May 2022.

# Use Cases beyond COVID-19

The Cardea ecosystem can be easily adapted to manage verification of other health data.

### Easily addressed with minimal changes to the existing code

- Health certification for school or work
  - Vaccines
  - TB testing
  - Annual physical examination
- Healthcare employee health requirements
  - Immunization proofs
  - TB testing
  - Recurring mandatory lab tests
  - Incident testing/followup (e.g., needle-stick incidents)
  - Drug testing

### Requiring some extra effort to address new use cases

- Communicable diseases
  - Mandated reporting (including the ability for a patient to share their status)
- Patient identity/profiles
  - Single source of patient demographics
- Release of data/consent
  - Need to define subsets for this use case:
    - Medical research
    - Sharing between care teams
    - Sharing across sectors
- Patient-doctor communications
- Prescription credentials (no more paper prescriptions)
  - Prescription management and tracking

### Requiring significant development effort

- Data collection from patients
  - Capturing genomic data (i.e. 23 and me, pregnancy) for storing and sharing with privacy layer
  - Support for customized health data
  - Capture of SDOH data and other screening tools
- Healthcare employee credentialing
  - Management of training credentials, certifications, CMEs
- Medical device and wearable data management
  - a1C monitoring
  - Sports bands
  - BP/cardiac readers
  - Digital Twins

# Conclusion

Cardea was developed as an open source, end-to-end ecosystem for managing COVID-19 and vaccine testing at the height of the pandemic and with the explicit goal of ensuring patient and traveler privacy. A series of successful trials in real-world conditions has proven the underlying open source technology, its ease of use and ability to integrate with existing systems, its ability to deliver seamless travel experiences where health status can be combined with other information, and, above all, its value to participants.

Continued development of the codebase by the Cardea Community Group at Linux Foundation Public Health has added important features, notably machine readable governance for managing information flows and offline functionality, as well as improved interoperability on Hyperledger Indy-based networks.

The result is that Cardea is more than a solution for managing COVID-19 testing; it is a way to manage any health-related process where critical and personal information needs to be shared and verified in a way that enables privacy and enhances security. It is able to meet the requirements of the 21st Century Cures Act and Europe's General Data Protection Regulation and in doing so, enable use cases that range from simple "proof of identity" to interoperating ecosystems encompassing multiple cloud services, organizations, and sectors, where data needs to be, and can be, shared in immediately actionable ways.

Open source, interoperable decentralized identity technology is the only viable way to manage both the challenges of the present — where entire health systems can be held at ransom through identity-based breaches — and the opportunities presented by a digital future where digital twins, smart hospitals, and spatial web applications will reshape how healthcare is managed and delivered.

To find out more about this award-winning technology, visit Cardea.app.

## The Cardea Project would like to thank

*SITA, for its vision and commitment to developing this technology and for its donation of the additional code to create the Cardea Project.*

*The government of Aruba, for its belief in the technology and its enthusiastic collaboration at every stage of development, and all the people in Aruba who helped to make the trials successful.*

*The engineering and operations teams at Indicio.*

*Linux Foundation Public Health, for its encouragement and support of the Cardea Project, and in particular Jim St. Clair and Jenny Wanger.*

*The Cardea Steering Committee and all the companies and individuals who have supported the Cardea Community Group.*

*The Cardea Community Group, for its members' input into this white paper, and in particular Sam Curren, Kirsty Gale, and Bill Henn.*

The technical diagrams in this paper were provided by Indicio.tech.

KuppingerCole provided the photograph from the 2022 European Identity and Cloud Conference.

CARDEA

LF PUBLIC HEALTH