Use Cases and Proposed Solutions for Verifiable Offline Credentials

a white paper from Rebooting the Web of Trust VII

by Michael Lodder (mike@sovrin.org), Samantha Mathews Chase (samantha@venn.agency), and Wolf McNally (wolf@wolfmcnally.com)

ABSTRACT

Self-Sovereign Identity is now a widely discussed topic, especially in the context of verifiable credentials/attributes attested by third-parties about an individual or entity that can be used as proof about them, such as a digital driver's license or passport. This has enabled new systems to be developed to address security and privacy issues.

In this paper we cover various scenarios where some or all parties have intermittent, unreliable, untrusted, insecure, or no network access, but require cryptographic verification (message protection and/or proofs). Furthermore, communications between the parties may be only via legacy voice channels. Applicable situations include marine, subterranean, remote expeditions, disaster areas, refugee camps, and high-security installations. This paper then recommends solutions for addressing offline deployments.

civic

Protocol Labs

sovrin

EXODUS
htc

interlinked
connecting intelligence

PTB VENTURES

TIERION

Microsoft

Sponsors for the Rebooting the Web of Trust VII Design Workshop

INTRODUCTION

All current solutions for verifiable credentials involve computer systems and networks. These systems perform all the complex cryptographic algorithms on a user's behalf, communicate with all involved parties, and are responsible for safeguarding the information. Developed countries such as those in Europe and North America have no problem using these platforms, but these solutions exclude parts of the word that do not have the same sophistication. For example, the vast majority of Africa does not have access to the internet; though many Africa locations do, connection speeds are incredibly slow. There are also possible scenarios where a computer system may not be available to individuals like refugees or missionaries. Finally, some high-security environments do not allow outside network connections. All of these situations must still permit users to prove attributes about themselves and for relying parties to validate that information. The goal of this paper is to provide and encourage consideration for situations where users operate in internet-hostile conditions.

PREMISE

For purposes presented in this paper, offline means any device or method that does not require an active internet connection. Offline credentials aim to fill this niche. The benefits to such a system include an individual being able to carry the credentials with them and being able to safeguard them with minimal devices while still allowing relying parties to cryptographically verify them. Offline credentials should not require the use of major computer systems or other powerful electronic devices, but may use them to implement pieces of the process. Otherwise, this voids the entire idea of offline credentials. Some scenarios require cryptographic material processing remain offline to prevent electronic compromises, alterations, or theft. Offline encryption is also not susceptible to a malware attack; and side channel attacks require cameras to record the person performing the encryption. Offline encryption systems have existed for centuries as ciphers. Only in the last few decades have offline cryptography algorithms become sophisticated enough to offer the more powerful features of modern cryptography like authenticated encryption and yield ciphertexts with uniform random distribution of character frequencies. Understanding simple encryption can also help to establish trust in systems.

Offline benefits encompass three categories: usability, deployability, and security. Each of these come with risks and limits that computers already handle or solve. One question is how this is different than having a physical credential like a drivers license today. A drivers license or passport cannot be cryptographically verified by hand, and all information must be shown when presented to a relying party. Offline credentials should also support selective disclosure where only the information that a credential holder allows is shared with a verifier. This paper discusses a toolkit that allows possible solutions to be implemented to create verifiable offline credentials along with their pros and cons.

Considerations

Before detailing offline use cases and solutions, it is necessary to cover the aspects of offline credentials that will be used to measure solutions before they are considered viable. Some offline ciphers require more sophistication and are more prone to mistakes but hard to break, while others may be simpler, with fewer mistakes possible, but not hard enough for an attacker with sufficient resources to break. Many of these terms are borrowed from The quest to replace passwords (particularly concerning usability and security) but adapted to offline credentials. Below is a list of those considered relevant for offline credential systems.

Parties

1. *Holder*: A person or entity that physically holds offline credentials.
2. *Issuer*: A person or entity that creates offline credentials for *Holders*.
3. *Verifier*: A person or entity that receives a presentation of credentials from *Holders* and verifies their truthfulness.

Usability

We define usability to mean the following:

1. *Memory Wise-Effortless*: Holders do not have to remember any secrets at all or possibly one secret for everything (e.g., pin to unlock offline device).
2. *Scalable*: Dozens of credentials shouldn't increase the burden for the user. "Scalable" is only from the user's perspective.
3. *Simple-to-carry*: Users carry a minimal additional physical object (electronic, mechanical key, piece of paper) that stores credential and cryptographic material and is powerful enough to perform the necessary proofs.
4. *Physically-Effortless*: Process does not require physical user effort beyond, say, performing a simple task like pressing a button or entering a passcode.
5. *Easy-to-learn*: Users who don't know the process can figure it out and learn it without too much trouble, and then easily recall how to use it.
6. *Efficient-to-Use*: The time the user must spend for each presentation is acceptably short. The time required for enrolling with a new issuer, although possibly longer than presentation to a verifier, is also reasonable.
7. *Infrequent-Errors*: The process a user must perform should succeed when done by an honest and legitimate person. In other words, the system isn't so hard to use or unreliable that genuine users are routinely rejected (as might occur when performing an authenticated encryption scheme by hand).
8. *Easy-Recovery-from-Loss*: Users can conveniently regain their credentials if lost or stolen. This combines other aspects like low latency before restored credentials; low user inconvenience (e.g., no requirement for physically standing in line); and assurance that recovery will be possible.
9. *Selective-disclosure*: Users can easily choose which attributes to present and withhold the rest.

Deployability

We define deployability to be the following:

1. *Accessible*: Holders are not prevented from using the system by disabilities or other physical (not cognitive) conditions.
2. *Negligible-Cost-per-User*: The total cost per Holder is negligible.
3. *System-compatible*: The process could be done by computers if needed.
4. *Non-Proprietary*: Anyone can implement or use the process for any purpose without having to pay royalties to anyone else. Relevant techniques are generally known, published openly, and not protected by patents or trade secrets.

This category is often the barrier for moving offline. Data collected offline could be lost before it has a chance to be saved to online resources.

Security

We define security to be the following:

1. *Resilient-to-Physical-Observation*: An attacker cannot impersonate a user after observing them present a credential. Attacks include any form of observation.
2. *Resilient-to-Targeted-Impersonation*: It is not possible for an attacker to impersonate a holder by exploiting knowledge of personal details without having their credentials.
3. *Resilient-to-Guessing*: Since offline credential presentations are done in person, relying parties can constrain guessing or detect an impersonator trying to guess.
4. *Resilient-to-Leaks-from-Other-Verifiers*: Nothing a verifier could possibly leak can help an attacker impersonate the user to another verifier.
5. *Resilient-to-Theft*: An attacker in possession of a Holder's credentials cannot use them for presentation to another party.
6. *No-Trusted-Third-Party*: The process does not rely on a trusted third party who could, upon being attacked or otherwise becoming untrustworthy, compromise a holder's security or privacy.
7. *Unlinkable*: Colluding verifiers cannot determine whether the same holder is presenting to both.

USE CASES

We begin by describing a scenario where offline credentials could be deployed, detailing their respective environments, and we conclude by illustrating how offline credentials are helpful.

Scenario: Marine, Subterranean

Underwater and underground operations present interesting conditions for internet connectivity.

Oil rigs, while stationary, are constructed many miles offshore and may have consistent connections to land-based internet or satellites but may not always be online. Offshore rescue operations like the US Coast Guard tend to have short gaps of time where connectivity does not exist, but usually return to land after a few minutes or a few hours. Boats can travel longer distances away from land and may be away weeks or months at a time. Since boats travel at the surface, they do have more internet accessibility than submarines. Submarines only connect to the outside world when at or near the surface. Information routing is critical for marine-based scenarios because often the data to be sent have directed destinations. An oil rig in international waters may not want to send highly sensitive data over the nearest countries internet cables but instead chooses to send it by human carriers. In [2], the authors describe certain weaknesses in underwater security communications and recommend “non-interactive data transmission schemes that ensure the underwater vehicles do not transmit additional messages for authentication and key establishment.”

Underground inhibits wireless signals and may require long cables to connect to operators. They might choose to report in or to sync to shore or surface at longer intervals, such as once a day.

Scenario: Remote Expeditions

There are remote locations like the Amazon jungle, Antarctica, or Mount Everest where it's possible to receive communications but which require expensive equipment to reach satellites or remote substations.

Scenario: Sensitive Compartmented Information Facility (SCIF)

SCIFs store high security information. The most restrictive security measures are implemented, such as fences, guards, faraday cages, concrete walls, and no internet connection. Personnel are required to undergo strict vetting for security clearances and utilize multiple factors for authentication. At designated intervals, updates to data are performed but methods vary, such as using USB thumb drives containing the newest data for transportation and synchronization.

Scenario: Epidemic Outbreaks, Public Health Events

Epidemics can occur with or without warning. When they happen, there are no guarantees whether internet connections are available, especially in developing countries like Africa. Tracking individuals who have received a clean bill of health and those infected becomes paramount. Affected areas are isolated via containment procedures, with restricted access to qualified medical professionals [5, 6]. Medical workers must be able to prove their qualifications and document their actions with or without internet access.

Scenario: Disaster Zones

Fraud also becomes prevalent as charitable donations are given to malicious parties. Internet connectivity can be sparse for emergency responders because important technology infrastructure is commonly destroyed or otherwise unavailable or disrupted. It is critical when disaster strikes to be able to identify those affected. Any rescued personnel’s infirmity or disability are necessary to know before administering medical services, in case of allergies or adverse side effects to medicine or treatment. Remote operated vehicles (ROVs) are commonly deployed to assist in rescue efforts. ROVs solely communicate with operators and over Radio Frequency and not the internet. ROVs could facilitate identification if provided with features to identify people.

Scenario: Missionary, Humanitarian Service, Refugee Camps

Religious missionaries or humanitarian aid workers are required to prove immunizations and to present visas and other documentation when entering countries for service [3, 4]. Their service could include documenting vital information on behalf of the local population and recording their activities while in internet dead zones. The information is eventually relocated to internet-connected areas but the time period is unknown. Many locals do not have a digital presence at all or have very little of a presence. Refugee camps are often required to immunize and document other medical facts and procedures on refugees.

Scenario: Resource Allocation and Management

Areas where essential resources like food, water, or medicine are scarce require careful allocation and management. Often such areas are poorly connected to the Internet and visited by personnel who need permission to access and distribute the resources.

Scenario: Delegative Democracy

Democracy is often implemented using “representative democracy” where citizens elect representatives in a particular district for a particular term, who “represent” their district as they please. This technique was developed when distances were large and communication was slow; it often concentrates power in the hands of elected representatives who are not very accountable to their constituents. By contrast, delegative democracy (also called “liquid democracy”) allows anyone to delegate their vote to anyone else. This also works well over long distances and in areas with poor connectivity, but requires more bookkeeping and auditing to hold delegates accountable. In particular, the ability to assign proxy credentials to delegates and remain anonymous, the ability for delegates to vote their proxies and be accountable, and the ability to prove the results of voting are all necessary.

SOLUTIONS

The scenarios and the variance in requirements are vast. However, many of the challenges can be solved in similar ways. We propose a toolkit that can be used to address most of the concerns, but implementation is left to developers. The toolkit includes methods for confidential identification, authentication, authorization, and auditing. It is the hope that implementers consider basing their solutions on this toolkit and the ideas proposed.

Sovrin is a blockchain for enabling secure privacy-preserving identity management, but any blockchain that facilitates identity management should suffice. The Sovrin ledger enables users to easily and securely manage identities by using verifiable credentials and zero-knowledge proofs. Organizations can securely issue credentials containing various attributes to personnel, which can be used to generate zero-knowledge proofs to relying parties. Issuers use Sovrin to indicate which credentials are valid without disclosing to whom they were issued, and Relying Parties can use the ledger to verify proofs from credential holders. In order to enable offline verification, Sovrin supports creating a proof of ledger state: a snapshot of the ledger at the latest moment in time. Devices can store state proofs and receive periodic updates as permitted, but would not require persistent internet connections. Updates can be performed according to best security and industry practices. Credentials can be stored in offline digital wallets, i.e. smart cards, USB keys, ROVs, and custom electronics. These devices can also support proof-request parsing, proof generation, and proof verification. All of the code to do this is open source in Hyperledger Indy-SDK. Sovrin supports many of the considerations stated previously and selective disclosure.

Offline still requires devices with sufficient computational power to perform complex math operations for issuing credentials and generating zero-knowledge proofs such as modular exponentiation and elliptic curve pairings but would not require anything else. It is recommended they also provide cryptographic primitives for encryption, digital signatures, and key exchanges to meet simplicity requirements for users.

Offline methods can securely transmit issued credential from Issuer to Holder over any medium: QR codes, audio codes, RF, downloaded to portable thumb drive, dead drops, etc. Offline transmissions can be secured by using modern key-exchange protocols like Diffie Hellman if supported but may also be done via pen-and-paper ciphers. If using pen-and-paper ciphers, care should be taken for the time it takes to perform encryption and to use authenticated encryption scheme. LC4 is a new cipher that supports authenticated encryption and nonces to prevent key leaks. The algorithm is easy to perform which makes it susceptible to some attacks but simple to learn. It supports up to 36 characters, keys are 36 characters long, and the algorithm claims strengths equal to 136 bits. LC4 has most of the considerations and allows the scheme to meet all the requirements. Zero-knowledge proofs do not need to be encrypted as they reveal no information,` unlike credentials. This could simplify designs and reduce power consumption if needed. It is recommended to use encryption in all message transmissions to provide security by default.

Some scenarios call for storing non-credential data that must be shared. IPFS is a distributed hash table (DHT) tool that allows offline state caching, modification and later resyncing. Again, any distributed data technology could work. This is helpful for collecting data about situations like radiation or atmosphere readings, soil samples, patient blood samples, immunizations, etc. This data can later be read and additional credentials can be issued, or it can be published to for global consumption like IXO to measure impact. IXO allows Sovrin credentials for authentically publishing data.

This toolkit of three technologies enable solutions to be easily implemented that meet scenario requirements. The blockchains have been created to be usable, deployable, and secure. We envision future technology built from this toolkit that solves the challenges from these scenarios in ways that consumers couldn't imagine. The authors are working on basic implementations that can be consumed for more complicated deployments.

To conclude, we explore how the marine use case could be solved using this toolkit:

Maritime communications happen via SATCOM (satellite communications). According to [7], security is rudimentary, weak, or non existent. The first problem is identification and authentication. Ship workers, sailors, and any electronic device can be given credentials that contain essential attributes. Credentials can be stored in smart cards for people, system memory for devices. Any device or system requiring authentication, proof of training, or proof of rank to make decisions can store state proofs from Sovrin. If the company demanded more security, Hyperledger Indy could be deployed on their own networks but would limit portability for cross-company collaboration. State proof updates can be done over SATCOM using authentication from credentials issued to the fleet while at port. Updates to devices on the boat can be done any viable method but for our example we state the boat has networked devices to a main computer which functions as the external gateway. Workers can access their work areas via smart cards that perform authentication with ZKPs. Audits for work completed can be recorded in the devices using IPFS. When the next resync window happens on this data can be sent securely to where it needs to go. This fulfills the usability requirements because workers just need to remember to carry their smart cards; it’s easy to use by simply sliding in a card or presenting the card via a proximity scanner. It's secure because it uses zero-knowledge proofs to verify identities, and deployable because the smart cards can be made from inexpensive hardware like a raspberry pi or embedded systems. Once at port, all the public IPFS data could be aggregated and uploaded for impact analysis to IXO for shipping companies to review later for improving their trade. Implementation does not change much or at all for offshore operations or submarines.

REFERENCES

[1] Pacific Life Research Center - http://www.plrc.org/docs/941005B.pdf

[2] Changsheng Wan, Vir Virander Phoha, Yuzhe Tang, Aiqun Hu, "Non-interactive Identity-Based Underwater Data Transmission With Anonymity and Zero Knowledge", Vehicular Technology IEEE Transactions on, vol. 67, no. 2, pp. 1726-1739, 2018.

[3] World Health Organization Humanitarian Requirements Document - http://www.who.int/health-cluster/countries/ethiopia/ethiopia-humanitarian-response-plan-2017.pdf

[4] Ethiopia Requirements Document - https://www.usaid.gov/sites/default/files/documents/1860/Ethiopia%20HRD%202016.pdf

[5] International Health Regulations and Epidemic Control - http://www.who.int/trade/distance_learning/gpgh/gpgh8/en/

[6] Control of communicable diseases and prevention of epidemics - http://www.searo.who.int/entity/emergencies/documents/who_control_of_communicable_disease.pdf?ua=1

[7] Hacking, tracking, stealing and sinking ships - https://www.pentestpartners.com/security-blog/hacking-tracking-stealing-and-sinking-ships/

ADDITIONAL CREDITS

Lead Author: Michael Lodder (mike@sovrin.org)

Authors: Samantha Mathews Chase (samantha@venn.agency) and Wolf McNally (wolf@wolfmcnally.com)

About Rebooting the Web of Trust

This paper was produced as part of the Rebooting the Web of Trust VII design workshop. On September 26th through 28th, 2018, over 40 tech visionaries came together in Mississauga, Ontario to talk about the future of decentralized trust on the internet with the goal of writing 3-5 white papers and specs. This is one of them.

Leadership Team: Christopher Allen, Joe Andrieu, Kim Hamilton Duffy, Manu Sporny, and Heather Vescent

Gold Sponsors: Civic, Protocol Labs, Sovrin, an anonymous donor

Silver Sponsors: HTC, Microsoft, PTB Ventures, Spherity, Tierion

Community Sponsors: Blockchain Commons, Learning Machine, Legendary Requirements, Purple Tornado, Veres One

Workshop Credits: Christopher Allen (Founder), Joe Andrieu (Producer and Facilitator), Shannon Appelcline (Editor-in-chief), and Claire Rumore (Graphical Recorder)

Thanks to our other contributors and sponsors!

What's Next?

The design workshop and this paper are just starting points for Rebooting the Web of Trust. If you have any comments, thoughts, or expansions on this paper, please post them to our GitHub issues page:

https://github.com/WebOfTrustInfo/rwot7/issues

The next Rebooting the Web of Trust design workshop is scheduled for the week of March 1st to March 3rd in Barcelona, Spain. If you'd like to be involved or would like to help sponsor the event, email:

rwot-leadership@googlegroups.com