Sentries, Confessionals, Vaults, and Envelopes

Daniel Hardman

2023-01-18

Abstract

Decentralized identity suffers from a recurring confusion: treating OpenID Connect, CHAPI, Decentralized Web Nodes, and DIDComm as rival solutions to one problem. This primer offers a medieval metaphor to keep them straight. OIDC is like sentries at a gate (real-time login and gatekeeping); CHAPI is like a confessional (structured, sandboxed credential sharing with a website); a Decentralized Web Node is like a vault (specialist, policy-governed storage of valuable data); and DIDComm is like an envelope (asynchronous, transport-agnostic messaging secured by DID control). Each tool has a sweet spot and a different relationship to time, boundaries, and trust. The broader argument is that interoperability is not served by forcing these tools to converge. The author distinguishes auto-interop — agreement among tools that share a goal — from xeno-interop, the harder work of drawing crisp boundaries between tools that solve genuinely different problems, and contends that decentralized identity needs more of the latter.

The decentralized identity / SSI space is troubled by persistent misunderstandings about the overlaps between important technologies: OpenID Connect (OIDC), CHAPI, Decentralized Web Nodes (DWNs) [and similar members of the secure storage family], and DIDComm.

Sometimes these are described as competing solutions. I think that's unhelpful. They do compete for our time and attention, I suppose, but they solve different problems for different reasons — and a clear understanding of these differences would go a long way to resolving contention in our community. So I'd like to offer an extended metaphor to make sense of it all. The metaphor will draw upon these 4 pictures:

But before I explain, let me offer two caveats.

1. *Metaphors are a two-edged sword.* They can demystify the unknown, and simplify the complex. However, they may also be applied too literally or too broadly, and they can imply color that is unintended. I recently used another metaphor about "grand unified theory" that triggered some reactions I didn't expect, so I'm wanting to be cautious. Please approach the following metaphor thoughtfully and tentatively; the intent is insight, not dogma or dissing.
2. *I am not an unbiased source.* I am one of the architects of DIDComm, and this shows in how I invest my time and energy. In other contexts, I do advocate DIDComm's virtues. However, I don't intend DIDComm advocacy in this article, only a helpful mental model, and I don't aspire (either here or elsewhere) to criticize the other technologies I enumerated. Make of that what you will.

Figure 1: image credit: Midjourney

TL;DR

OIDC ≅ (is analogous to) sentries. CHAPI ≅ confessionals. DWN ≅ vaults. DIDComm ≅ envelopes. The first two technologies are login-oriented. DWNs are data-oriented. DIDComm is communication-oriented. Each tool is useful, and each has a trust problem that it solves particularly well. A smart sovereign will define their trust goals carefully, and in a given circumstance, use whichever tool best achieves them.

A lot of confusion and dissonance in SSI-land is caused by the mistaken belief that interop requires the unification of these tools, or a "win" for one of these tools at the expense of the others. Our community also has unspoken disagreements about the relative value of different sovereign goals. This influences perceptions of the relative usefulness or attractiveness of the tools, which causes more dissonance. We might get better results if we used a different mental model. Hence, the metaphor.

The Basic Metaphor

Imagine a world like the kingdoms and castles of medieval Europe. Alice is a queen who lives in a castle; Bob is a king who lives in a walled city across the river. Other kingdoms exist, and the land between is full of vagabonds and other dangerous characters. We can map the issues we wrestle in SSI onto trust challenges in this landscape, and the technologies of SSI onto medieval tools.

One important trust mechanism available to Alice and Bob is *sentries*. These are soldiers posted at the gate of a castle or a walled city, who vet any party seeking to enter.

Another mechanism is *confessionals*. These allow people to enter a carefully curated context where sensitive personal information can be disclosed under rules enforced by ecclesiastical authorities.

A third mechanism is *vaults*. In Alice and Bob's world, bulk management of valuables means treasure chests behind strong doors with armed guards. Since Alice doesn't want to run down to the vault with the keys on her belt every time she makes a payment, she creates written policies for the treasurer. A treasurer's guild has developed special skills in this space, and has a reputation for following such instructions to the letter. The existence of this specialized guild allows another valuable service: bulk movement of treasure, under armed guard, from one place to another.

A fourth trust mechanism is *envelopes*, including official seals, signet rings, wax imprints, ciphers, and related conventions for secure communication.

Sentries

Sentries are a credential-oriented authentication and/or authorization mechanism. The classic middle-of-the-night challenge "who goes there?" is pure authentication. Sometimes sentries make attribute-based decisions instead, or in addition ("I don't know this courier at all, but they're an official emissary from the neighboring kingdom, so I'll convey them under guard to the keep.")

Sentries are boundary- and session-oriented. They are charged with gatekeeping access, and they grant visitors a pass (session token) that's assumed to be valid as long as they are inside the castle or city walls.

Figure 2: image credit: Midjourney

Sentries always operate in real-time. Nobody expects a sentry to call "who goes there?" and wait hours for a response.

Sentries are good at denying entry to potentially hostile outside forces. They protect institutions with a gated boundary. Sentries aren't good at implementing complex treasury management, they don't help unburden the soul, and they're not a great way to courier confidential romantic letters between Alice and her one true love. This is not a value judgment on sentries, just a clear-eyed summary about fitness-for-purpose. A rational sovereign might use sentries as one component of a larger strategy for trust, but probably wouldn't use sentries alone.

OIDC has many of the same characteristics as sentries. It is fundamentally a login mechanism, charged with gatekeeping at the perimeter, that grants session-based access. It requires a challenge and a response that are credential-oriented.

If we're relating Alice's sentries to the original OIDC (what's widely deployed today), then the sentries are third parties, independent from the end user or the relying party. We might think of them as a guild that Alice hires and trusts to run sentry duty for her. The recent innovation of Self-Issued OpenID Providers (SIOP) is a clever way to change the logistics. Sentries now work for the end user, but Alice can trust them because they present evidence that is cryptographically unforgeable. When SIOP is broadly supported, it will add some nifty flexibility. However, I don't think this changes the metaphor much.

Confessionals

Confessionals are a physically constrained space in which it is safe to do one-way sharing of confidential personal information with a church official (an institution). They are governed by a well understood set of conventions and some ethical expectations about privacy and (lack of) record-keeping.

Confessionals are session-based, meaning that you must create and sustain certain conditions (walk into the church, open a door, sit down, stay there) to use them. They are also protocol-oriented. There is a recipe for interaction between priest and visitor. It has a beginning, structured turn-taking, and an end. The goals of the interaction are crisp.

Like sentries, confessionals are a real-time mechanism.

Confessionals are free of cost, and since the church is pervasive in Alice and Bob's medieval culture, they are also ubiquitous and staffed round the clock.

Like sentries, confessionals have a sweet spot. They are good at safe disclosure of a certain type of information to a certain type of listener. Using confessionals as a way to gain access to the treasury might be theoretically possible, but is probably suboptimal. Exchanging ciphered love letters inside a confessional also feels a bit clumsy.

CHAPI is a protocol for sharing credentials (typically, for login), specifically with a website. CHAPI interactions take place within the carefully sandboxed safe space of a browser session. There is a protocol. Over and above the protocol, browser vendors have written and unwritten rules about how their implementation of the API works (e.g., what other parts of the browser's memory sees what the API handles).

Because browsers are ubiquitous and free, nearly any member of the public can use CHAPI. The websites that sit on the other side of the "confessional" from a private individual are staffed

Figure 3: image credit: Midjourney

around the clock. Just like confessionals need a church, CHAPI needs a browser.

One flaw in the CHAPI ≅ confessional parallel is that confessionals are passive, whereas CHAPI is active. Priests neither seek nor react overtly to confessions. On the other hand, CHAPI helps websites challenge a user for credentials, and the website reacts by authorizing access (or not). This makes CHAPI much more sentry-like than the analogy suggests. We could stretch the metaphor by imagining that confessionals are like heaven's gatekeepers, or we could say that in Alice's world, visitors access the worldly castle either through the front gate or by seeking priestly approval at the church. Either way, the limits of our metaphor are showing. Keep this in mind as we proceed.

Vaults

Treasure is a big trust issue in our fictional Alice and Bob world, and vaults specialize in treasure management.

Figure 4: image credit: Midjourney

Vaults are passive and stationary. Those who want something out of the vault must approach the boundary and present proof that they are entitled. The protected resource is then released.

Vaults are oriented toward heavy-duty, bulk operations. They have shelves and chests. They are big enough to store a hoard, and they allow groups to walk inside when wrestling chests. Multiple people might consume parts of the space of a shared vault. Vaults are armored and strong.

With respect to time, vaults have the interesting property that they decouple the timing of the owner and the timing of the accessor. Each interacts in real-time, but those real-times don't have to align because the vault is operated by a third party.

Despite these wonderful qualities, not all trust issues have a natural affinity for vaults. Vaults do not replace sentries or confessionals. Vaults do not even store all precious material. Alice might wear a valuable ring on her hand. She might value her daughter as a treasure, but she wouldn't store her daughter in a vault. Alice might prefer to send, receive, and store her love letters without running down to the dark reaches of her castle, and without the involvement of a treasurer's guild.

If the treasure of decentralized identity is data, then DWNs (or similar technologies like SOLID pods, encrypted data vaults...) are the SSI analog to vaults. These tools are all about managing valuable data, and they are offered as a service by specialized providers (a guild).

DWNs are server-based, and they expose an API to clients. Data owners insert their data treasure into the vault and configure policies about how / when / under what conditions it may be released. Clients call the API to gain access.

A limit to the vault ≅ DWN parallel is that physical and digital management have different efficiencies due to machine speeds and automation. Data replication is a bit like armored caravans operated by the treasurer's guild, but it's safer and far less error-prone. Also, encryption can be used with DWN to partially protect treasure from embezzlement of a corrupt treasurer's guild. Despite these issues, I think the comparison remains instructive.

Envelopes

When Alice and Bob want to communicate securely, they are likely to use a combination of tools that focus specifically on messages. Collectively, I'll call these tools "envelopes", but they are actually a bit more complex.

Alice writes a letter on parchment, signs it, and places it in a protective envelope. She then adds a tamper-evident wax seal to make it hard for prying eyes to eavesdrop. If she wants additional security, she may write using a special code that requires decoding by the recipient — and this may require the recipient to have decoding helps that Alice shares in advance. Alice may also employ a system of couriers, dead drops, and proxies to arrange confidential delivery.

Notice that this mechanism is multimodal (could move over land or sea or carrier pigeon) and one-way. Just because Alice sends a letter by courier X on road Y does not mean Bob has to send a reply back by the same courier on the same road.

Envelopes are also fully asynchronous. Neither party has to be present in real-time. This is one degree looser in its temporal coupling than vaults, and two degrees looser than sentries or confessionals. That has pros and cons.

Just like all the other tools we've discussed, envelopes have a sweet spot. They are great for directed interpersonal communication between pairs or small groups of peers — and because basic communication is foundational to fancier interactions, they can negotiate treaties or build a romance. Envelopes can also carry public announcements. However, sentries might demand to see a face rather than a letter before they open a gate, priests give better counsel when they

Figure 5: image credit: Midjourney

hear a tone of voice, and Alice would never store the full 90-volume history of her kingdom in her vault by enclosing each page in an individual envelope.

DIDComm is the toolset that corresponds to envelopes. It is a message-oriented communication methodology that derives its security from DID control. Digital signatures and encryption algorithms are the analog to wax seals and primitive ciphers. Mediators (a bit like SMTP mail transfer agents) are the analog to couriers, dead drops, and routing instructions.

Like letters, DIDComm messages are multimodal (the common term in DIDComm land is "transport agnostic"). This makes DIDComm unusual among SSI technologies in that HTTP and web tech is optional rather than required. Other supported transports include QR codes, Bluetooth, NFC, LoRa, email, sneakernet, and so forth. DIDComm is also less session-oriented and less client-server oriented than the other tools we've been discussing.

Like email, DIDComm supports totally asynchronous interactions.

One persistent confusion in decentralized identity is to equate DIDComm with credential exchange protocols that build on DIDComm (e.g., Aries credential exchange; WACI). Envelopes can carry trust deeds (very well), but they are far more general purpose; DIDComm can carry credentials (very well), but its intended use is broader. I spent a year writing digital cash protocols with DIDcomm that have little in common with credentials.

Stepping Back

Now that the metaphor is drawn in broad strokes, I'd like to do a little bit of reasoning with it. As I said earlier, my purpose here is to bring some clarity that's been lacking; I'll save advocacy for a different time.

1. All of these tools have legitimate uses. Vaults (for example) are not "bad" because sentries are "good."
2. A mature ecosystem will probably have all of them.
3. There is, *and always will be*, some overlap. It will always be possible to put confessions in an envelope, or to store love letters in a vault. This is not a bad thing, and it is not evidence of maverick intransigence, stupidity, or wasteful reinvention. What IS a bad thing is to misunderstand, misdescribe, undervalue, or mistarget the sweet spot of each.
4. Interoperability is not best served by making sentries more like envelopes, or confessionals more like vaults (for example). It's also not best served by pressuring the treasurer's guild and the sentries (for example) to find common ground. *These mechanisms have different goals, and each is valid.*

In SSI, I believe that interoperability is best served by subdividing the problem space, based on some careful thought about clusters of worldviews with high and low cohesion. Where a cluster has high cohesion, we should encourage clumping; where cohesion is low, we should create boundaries. This allows experts to work inside harmonious spaces that demand their skills and excite their passions, and it creates two *different* kinds of alignment efforts: *auto-interop* ("auto" from the Greek root meaning "self" or "directed within") inside a worldview that has high internal cohesion, and *xeno-interop* ("xeno" from the Greek root meaning "foreign") that bridges worldview affinity boundaries.

These two types of interop effort require different strategies. Articulating and accepting the distinction will help us improve our dialog because it changes the questions we ask.

Auto-interop is a good goal when parties agree on goals and priorities. It is driven by the question: *"What is our consensus about what must be common in an optimal recipe?"* This is a standards-body question. Every participant inside a worldview affinity cluster is making tools with similar features (or tools that constitute part of a single logical feature set). They compete by serving different customers, making different price/performance tradeoffs, or being more clever or better partnered or whatever. Here, interop is trying to guarantee that customers can replace one vendor or one stack with another that's reasonably equivalent. Vendors of car parts want auto-interop because it lets them create many products from a single process, and sell nuts and radiators and fan belts to many buyers. Buyers want it because it gives them many vendor options.

Xeno-interop happens when parties DON'T agree on goals and priorities, because they are attacking different problems. There, interop is driven by the question: *"How do we recognize and delineate boundaries between our problem domain and a neighbor's as crisply as possible, so we are differentiated and collaborative across the boundary rather than creating competitive chaos?"* This is NOT a standards-body question, and answering it feels more like an art than a science. If we achieve xeno-interop, dissimilar things can combine to accomplish a whole that neither does well on its own.

I suggest that the distinctions that I laid out in my extended metaphor above call for xeno-interop, NOT auto-interop, and I suggest that we haven't spent enough energy answering the xeno question about where to draw boundaries. Instead, we see a lot of groups that are trying to totalize everything, and a lot of finger wagging about the need to come together without a deep appreciation for the value of the diversity.

In the medieval world of Alice and Bob, it makes no sense to push for a standard that covers sentries, confessionals, vaults, and envelopes — and in SSI, it makes no sense to unify OIDC, CHAPI, DWN, and DIDComm. It makes no sense to frame them as competitors in a bake-off, either. They're mostly differentiated and complimentary. That does not mean they necessarily deserve our equal love, but it does mean they're validly independent efforts.

Some people in our space love login and / or authZ of people to institutions, and think that the *most* important place to focus is on this problem. They're not crazy; many institutions that pay for SSI start here. This is the sweet spot for OIDC and CHAPI.

Other people think it's *mostly* about data, and they're pursuing DWN. They're not crazy either; secure online storage is a recognized product category with clear demand and many use cases.

Other people think it's *mostly* about composable interactions based on communication primitives, and they're pursuing DIDComm. Since I'm in this camp, I won't claim I'm not crazy. I'll just say I hope I'm not. :-)

Notice that I used the hedge word "*most*[*ly*]" in the last three paragraphs. The claim about importance implicit in this word might be the source of some controversy, but its imprecision is a way to give one another grace and space. If we can agree that there is at least some room for the other approaches, and if we can agree that "how much room, exactly?" is not yet known or knowable, then we can co-exist cheerfully and let the market teach us. If we do that, and if we work on clarifying goals and boundaries instead of totalizing the problem domain, I'm confident that Alice and Bob will end up being happy, safe, and empowered sovereigns.