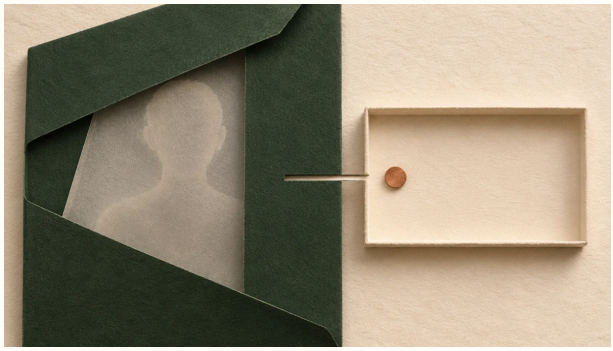
SAFWAN ALMODHAYAN



OPINION AND ANALYSIS

Editorial date: October 8, 2026

Prove Who You Are Without Telling Me Who You Are

When a Smaller Fact Is Enough

In California, a digital driver's licence can tell a shopkeeper less than its plastic counterpart and be more useful for doing so. Its holder can present an electronic verifier with proof that they have reached a required age; in some transactions, the party checking the credential may receive confirmation that the threshold has been met, while others may require more identifying information. The California Department of Motor Vehicles also makes clear that the digital licence is still accepted only at selected places and services. The surprise is not that the licence has moved from a wallet to a phone. It is that a question we used to answer by exposing a name, photograph, address and date of birth can be answered with a much smaller fact: this person is old enough.1

The change may look like a way to shrink a large document. In fact, it redraws the boundary of knowledge. What does the shopkeeper actually need to know? Not the customer's name, birthday, licence number or an address that has nothing to do with the purchase. The shopkeeper needs a decision with an acceptable degree of legal and technical confidence: has this person reached the age set by the rule? The old design handed an employee an almost complete personal file, then relied on them to ignore most of it. The new design tries to make what the system can see equal to what it has been authorised to know. The point is no longer only to protect data after collecting it. It is to ask why that data entered the room at all.

The word “identity” changes meaning here. In contemporary digital systems, it performs several jobs that habit has bundled under one name. Someone opening a bank account subject to Know Your Customer (KYC) requirements may have to prove and verify their legal name. A candidate entering an official examination may need to establish that they are the particular person who registered. But someone proving that they attend a recognised university, live within a district whose residents qualify for a service, or completed mandatory training before entering a facility does not necessarily need to disclose a comprehensive copy of their life. The distinction is not always between identity and anonymity. It is between knowing a person and knowing one relevant attribute about them.

From Card to Credential

In July 2025, the European Commission introduced a technical blueprint for age verification that would let countries and digital services build applications around proof of age eligibility without exposing unnecessary identifying information. Testing and adaptation with countries, platforms and users began that month. A second version followed in October, adding ways to rely on passports and identity cards; the Commission said the solution included zero-knowledge proof techniques. By April 2026, it said the application was technically ready and would be made available to citizens later. It named Denmark, France, Greece, Italy and Spain among the countries adapting it. These are facts about development, testing and a deployment plan. They are not evidence that every European was already using the system when this essay was written.2

The European blueprint points to an ambition deeper than improving a login screen. To issue an age credential, an issuer relies on a trusted source such as a national identity, passport or another approved means of verification. The design then aims to separate use of the proof from the organisation that issued it, so the issuer is not notified of every site its holder visits. The digital service requests proof of the threshold and receives a result with which to permit or refuse entry. Two moments that institutions habitually merge have been separated: establishing the trustworthiness of the source, and using a defined fraction of what it knows. That may sound like a fine technical distinction. It can also prevent ordinary acts of proving eligibility from becoming a central ledger of a person’s life.2

In May 2025, the World Wide Web Consortium (W3C) adopted version 2.0 of the Verifiable Credentials Data Model. Its basic idea is that a trusted organisation can issue claims about a person or entity in a form that allows someone else to verify the integrity of the credential and its relationship to its source. A university might attest to a degree, a professional body to a licence, or a government authority to an attribute in its records. The model names three parties: the issuer, the holder and the verifier. This structure is not a single global database. It allows a trusted claim to travel between organisations with different systems, provided they agree on the meaning of the credential and the authority of whoever issued it.3

But “verifiable” is not a synonym for “true.” The standard itself warns that being able to verify a credential does not necessarily prove that every statement inside it corresponds to reality. A digital signature can reveal tampering and connect the data to its issuer. It cannot travel back through an academic year to confirm that a graduate attended lectures, or summon a doctor to evaluate the examination they signed. If the issuer’s database is wrong, the wrong fact can enter the outside world in a cryptographically immaculate envelope. The system must therefore begin with questions about an institution and its procedures, not merely with a choice of signature algorithm.3

Disclosure Is Not Unlinkability

Imagine a professional certificate containing a name, date of birth, membership number, accreditation date and specialism. A work platform wants to verify that its holder has a valid licence in a particular field. The first method, and the most familiar to an administrative mind, is to send the whole certificate. The second is to issue a new version containing less information. The third is to let the person hold a composite credential and disclose only what the transaction requests. That third method is selective disclosure. If the person can prove that the licence is valid, or that a value lies above a threshold, without revealing the original value from which the proof was constructed, we approach another family of tools: zero-knowledge proofs. These are not interchangeable marketing terms. Each has different mathematical and security limits.

In November 2025, the Internet Engineering Task Force published RFC 9901, a standard for selective disclosure of signed JSON Web Tokens known as SD-JWT. It allows an issuer to create a credential containing cryptographic digests of fields, after which the holder chooses which fields to reveal to a verifier. In a simple example, the holder can disclose the attribute “student” while withholding an address and phone number, provided the fields were designed for separate disclosure. This reduces the information reaching the other party. Neither its name nor the mere presence of cryptography, however, proves that one presentation of the credential cannot be linked to another.4

This is where one of the market’s most insidious confusions appears. An application may claim to protect privacy because it does not show a user’s birthday, yet transmit the same signed token or another stable fingerprint on every occasion. Two cooperating sites may then discover—perhaps without learning a name—that the visitor who proved their age here is the same person who proved it there. RFC 9901 states this limitation explicitly: disclosure mechanisms based on salted hashes, including SD-JWT, cannot ensure separation between issuer and verifier against a colluding verifier that sends the original credential back to its issuer. We hid the field, but may have retained an invisible serial number from which to build a new biography of the visitor.4

Researchers therefore distinguish three properties that a designer should not allow to blur together. The first is disclosure minimisation: the verifier learns no more than it needs. The second is unlinkability: two verifiers cannot infer from the cryptographic material alone that two presentations came from the same credential. The third is holder binding: another person cannot copy the credential and present it as their own. Strengthening the third may weaken the second if every verifier sees the same persistent key. A system may stop credential sharing while tracking every use. This is not an impossible contradiction. It is a set of design conditions that must be tested together, rather than a reason to declare victory over one and silence discussion of the others.

What Does the Proof Actually Prove?

The most intriguing mathematical solution begins not with a person's name, but with a proposition that can be computed. Let a hidden value represent the person's date of birth, while the date of the transaction is known. The task is to convince the verifier that the difference between those dates meets the permitted age. In a simplified form, the statement might be: there exists a trusted original value for which "age is at least eighteen on the day of verification" is true. If the proof is sound, the verifier does not need the value itself. One crucial condition is hiding in the statement: the system must establish that the secret date came from a trusted source, rather than from a date typed by the user moments earlier. Otherwise, we have a beautiful mathematical proof of a fabricated fact.

In theory, a zero-knowledge proof combines three requirements. A verifier accepts a valid proof when the statement is true. Someone without a valid witness cannot persuade the system of a false statement except with the small probability permitted by the security assumptions. And the proof reveals nothing about the hidden witness beyond what the truth of the proposition itself discloses. Researchers call these properties completeness, soundness and zero knowledge. The third does not mean that the sentence "I am over eighteen" reveals nothing about the person. It means, within the relevant mathematical model, that the proof leaks no more of the witness. It does not by itself conceal a face before a camera, a network address or a browser history.[5]

Nor can every signed certificate automatically produce any proof we might wish to ask of it. If the issuer has already attested that the holder is "over eighteen," presenting that signed attribute may suffice without revealing a birth date. If the issuer has signed only the date of birth, proving a numerical comparison without exposing it depends on the cryptographic scheme, the encoding, and the way the signature is related to the attribute. A system may need a credential derived by the issuer or a proof protocol suited to this calculation. We cannot place a ZKP badge on a screen and assume that verification of any condition has become free

or guaranteed. The engineering problem is not merely to hide a number, but to establish the source of the number, its relationship to the condition, and the time at which that condition applies.

In September 2026, W3C published a Candidate Recommendation Draft for BBS cryptosuites used with verifiable credentials. These techniques allow derived cryptographic proofs to reveal selected attributes while preserving unlinkability between derived proofs when used according to their defined properties. The September document was not a final Recommendation; it was a draft at the implementation-testing stage. That distinction matters in a market fond of presenting emerging technology as settled consensus. Unlinkable cryptographic material also cannot guarantee an unlinkable application if a user supplies unique information alongside it, a site writes a recurring identifier to the browser, or another link is retained outside the protocol.6

When Context Reveals What Cryptography Concealed

The picture changes again when we move from proof to person. A graduate from a small university might disclose a rare specialism and a particular graduating class, and those two true facts may be enough to infer their name even if the proof carries no direct identifier. An employee on a team of three might establish their eligibility to use a sensitive system; their role and the timing may reveal who used it. No algorithm has been broken. The interaction has disclosed a context narrow enough to turn the unknown into the known. Privacy should not be measured only by the number of blank boxes on a digital card, but by the size of the crowd in which its holder can still hide after disclosure.

Another, quieter layer of knowledge remains. A site may know neither a name nor a birth date, yet know the network address, device type, pattern of movement and time of visit. It may use profiling or commercial analytics that connect this login attempt to earlier activity. A strong protocol can stop credentials being linked cryptographically across institutions. It cannot prevent an institution from tracking a visitor by other means. The reverse is true as well: network safeguards and data-retention practices may be excellent even when the credential technology does not provide mathematical unlinkability. Privacy is not a certificate bestowed by a software library. It is a property to be measured across the service as a whole.

Revocation, Possession and the Live-Status Paradox

The philosophy of the European blueprint becomes clearer when we ask about the issuer. If a platform must contact the issuer of an age credential every time someone visits a site, that issuer may become a witness to the holder’s habits. Offline verification, or careful use of public keys and status lists, can reduce this problem. But a credential may be revoked or expire, its

holder may lose a phone, or its keys may be compromised. A verifier then needs to know whether the certificate remains valid. The US National Institute of Standards and Technology's 2025 revision requires attribute bundles to support revocation, recommends limited validity periods, and warns that a revocation-status service should not tell an issuer that a particular credential was presented to a particular service.[7]

This produces what we might call the live-status paradox. We want a certificate that can be verified without contacting its issuer, lest verification become surveillance. At the same time, we want every verifier to know immediately when the certificate has been cancelled because a device was stolen, eligibility changed, or a licence was revoked. There is no magic design that combines zero communication, zero revocation delay and no state held anywhere. There are architectures and trade-offs: short validity periods, public status lists with sufficiently coarse granularity, suitable non-revocation proofs, recovery and reissuance procedures, and risk tolerances that differ by service. When a company sells a credential that "never contacts the issuer," we should ask how it knows the credential was not revoked yesterday.

The next practical question is who is standing in front of the phone. If one person obtains proof that they are an adult and sends it to a younger friend, the privacy of the protocol does not rescue its purpose. Some designs therefore bind possession to a private key held in protected hardware and require a wallet to produce a fresh signature in response to a verifier's unique challenge. A random nonce reduces the chance that an old message can be replayed, while binding the response to a particular verifier stops a proof sent to a shop from being presented again to another platform. These measures constrain impersonation and replay; without additional methods, they do not establish that the person holding the device is the same person who completed the initial enrolment.[7][8]

California matters again here, not because it has achieved an ideal state, but because it puts verification back in its proper place. The DMV stresses that its mobile driver's licence is not simply a picture of the card and that presentation occurs through electronic verification; it recommends that users not merely display the licence screen for someone to inspect by eye. The challenge for shops and service operators is to provide compatible readers, understand what is requested and received, and train for a dead phone or an unavailable scanner. What good is a privacy-preserving credential inside a wallet if, when the reader fails, an employee requires its holder to hand over a complete image of the identity document?[1]

From the Laboratory to the Institution

We can imagine broader applications, provided they remain design hypotheses rather than achievements claimed to exist. A website advertising a student-only internship might request proof of current enrolment at a recognised institution. A research centre handling sensitive data might ask for proof of recent ethics training without copying certificate numbers. A city

might grant a benefit to residents of an administrative area without requesting an exact home address if eligibility turns only on whether the address falls within that area. A professional association might confirm that a practitioner's licence is valid without revealing a separate document number. In every case, success means issuing an inspectable attribute suited to a particular decision, not copying a document and redacting what we dislike after receiving it.

Systems handling money quickly expose the limits of the slogan. When someone opens a bank account in a jurisdiction that requires customer due diligence, the bank cannot replace its legal obligation to know its customer with an abstract sentence saying "this person is over eighteen." The Financial Action Task Force's March 2020 guidance on digital identity says that use of digital systems for customer due diligence depends on the assurance level of the identity system and the risks of the financial relationship. Digital verification is a tool for more precise compliance, not an escape from due diligence. Institutions can still minimise attributes irrelevant to the transaction, or verify a subsidiary eligibility requirement through another party without exposing all details again. Privacy does not mean ignoring a duty. It means knowing where duty ends and appetite for collection begins.9

This changes the structure of competition between companies. A platform that routinely collected every user's date of birth, address and identity copy in the name of "easy verification" was in fact building a vast and dangerous information asset it might not need. Every additional record must be secured, given a retention period, and explained; depending on the applicable law, authorities or individuals may need to be notified when it leaks. Information changes from marketing fuel into a liability that follows the organisation for years. An organisation accepting a trustworthy eligibility proof may reduce its data attack surface, but it still must secure verification evidence, session restrictions and fraud events it retains. The economics are not merely about deleting a database. They are about choosing an institution's responsibility for knowledge before that responsibility becomes a security incident.

The European Union's rules acquire practical importance from this angle. The November 2024 implementing regulation on European digital identity wallets speaks of transferring the minimum necessary data to authorised relying parties and separating transactions in ways that limit linkability. It places attribute requests and verification of requesting parties within a defined framework. Technology alone does not persuade companies to stop requesting a name. Registration and governance make it possible to ask whether the organisation has a legitimate reason to request it. A technical standard supplies the means; an institutional rule decides when that means may be used and within what limits.10

One social constraint remains beyond the reach of a stronger signature. What happens to someone without a recent phone? To a person who loses a device and its private key? To someone unable to obtain a source document because the underlying register contains conflicting data? To an employee who refuses to hand a personal phone to an employer? A

system built for “privacy by design” can become a new gate of exclusion if it is the only path to a public service. It needs recognised alternatives, secure recovery that does not let a fraudster seize a credential by deceiving support staff, a way to correct issuer errors, and clear responsibility when a valid proof is refused. These are not kindly human additions to a cryptography project. They are conditions for the system to remain sound once it leaves the laboratory.

An institution can be understood as a publisher of questions, rather than of facts about people. In every interaction, its system asks: may this person do this thing now? When the question is precise, we can engineer a limited answer. When it is vague, the list of requested data expands as insurance. The first step in building a privacy-respecting service is therefore not choosing a wallet or contracting with a proof provider, but decomposing the decision itself. Which attribute changes the outcome? Who has authority to attest to it? How long does it remain valid? Is it enough to know the attribute exists, or must we know the specific holder to prevent one person claiming a unique entitlement more than once?

That final question reveals a challenge in designing one-time benefits and entitlements. An agency may want to confirm that an applicant qualifies for assistance only once without creating a general file that links their use of other services. Merely proving “eligibility” does not stop a person presenting the attribute ten times. A scope-specific pseudonym or single-use token can be created for the programme, provided it does not become a universal tracking key. Some cryptographic schemes permit forms of domain-specific pseudonymity, and the BBS draft discusses ways to derive them. Designing this for a public service requires testing prevention of duplicate payment against unlinkability outside the programme, with a right to review errors. Secrecy without enforceability can destroy public confidence; enforcement without limits can destroy privacy.6

The Disclosure Budget

To keep these principles from becoming public-relations copy, I propose an initial engineering exercise that I call a disclosure budget. It is not an existing regulatory standard or a published experimental result. It is a design framework that can be tested. The institution starts with one sentence stating the decision to be made. For every attribute, it records why it is necessary, its source, who sees the underlying value, who sees the proof, how long it is kept, and whether uses can be linked. It then runs two adversarial exercises: combine logs from hypothetical verifiers and search for repeated links; then assume one verifier cooperates with the issuer and ask whether a particular credential holder can be identified from the presentation trace. Finally, test revocation, loss of the phone, repeated requests and network failure. A lower field count is not success if another persistent trace makes tracking possible.

The central metric cannot be only “what share of users completed verification?” We should measure the share of correct decisions reached without excess disclosure, valid proofs refused, false claims accepted, time required for revocation to propagate, cases forced to fall back to disclosure of a full document, and linkability between verifiers under a defined threat model. Samples must include modest devices and users with different ages and skill levels. A system can pass its cryptography tests and fail as a service. An extremely easy interface can also grant a fraudulent credential confidence faster than paper ever did.

As these systems mature, a new market will appear, though not one for selling new “identities.” It will include specialist issuers of attributes, wallets that protect holders’ keys, verifiers that administer attribute-request policies, trust bodies that manage standards and the keys of approved institutions, and tools for privacy assessment and fraud-resistance testing. The value worth building lies not in knowing more about people, but in being able to rely on less. This is why launching a glossy application labelled “decentralised identity” is not enough. The wallet may sit on a user’s phone while every decision remains subject to an allowlist controlled by a single operator, or while the issuer is signalled on every use. Decentralising storage is insufficient if control of the decision remains centralised.

There is also a philosophical temptation to resist. Some will say that the system most respectful of a person is one that knows absolutely nothing about them. A responsible service sometimes needs a documented fact and an identifiable party who can be held accountable. The goal is not to turn every citizen into a ghost, but to make an institution’s knowledge proportionate to the entitlement or risk it manages. At some doors, a name is a legitimate requirement. At others, knowing it is an act of epistemic trespass with no useful function. Nor should zero-knowledge proofs be made to promise resolution of legal and social conflict. Their strength is that they expand what can be proved mathematically without exposing the whole witness. What ought to be proved, to whom, and who has the right to demand it are questions cryptography cannot answer alone.

Institutional Humility

California’s digital licence tells a story of partial movement from an exposed card to the requested attribute. The European effort tells another: institutions trying to separate issuance of trust from surveillance of its use, through development and deployment stages whose completeness should not be overstated. SD-JWT shows that selective disclosure does not necessarily make a credential untraceable; BBS work points to deeper protection for the presentation trace itself, while requiring us to distinguish a draft from a final standard. NIST reminds us that revocation, key protection and error handling are not footnotes to identity, but part of its honesty. These examples are not competing for the title of “most cryptographic.” They expose layers of a problem we concealed for too long behind a photocopied identity card.

These technologies may lead us towards a rare virtue in systems: institutional humility. A service asks for no more than will qualify it to make a decision, and then stops. The greater achievement would be neither that a person can hide their name nor that a complicated proof succeeds before a fast processor. It would be that an institution learns the difference between its right to verify and its curiosity to know. We spent decades teaching machines to recognise us everywhere. It is time to teach them a harder skill: believe us on one particular matter, then close the file and leave the rest of our lives alone.

Sources and references

[1] California DMV, Mobile Driver's License FAQs — Current official page; reviewed 9 November 2026.

[2] European Commission, Blueprint for an age verification solution — First version July 2025; second version October 2025; technical readiness announced April 2026; page updated 22 September 2026.

[3] W3C, Verifiable Credentials Data Model v2.0 — W3C Recommendation, 15 May 2025.

[4] IETF, RFC 9901, Selective Disclosure for JSON Web Tokens — Published RFC, November 2025; see §§4 and 10.1.

[5] Goldwasser, Micali and Rackoff, "The Knowledge Complexity of Interactive Proof Systems" — *SIAM Journal on Computing* 18(1), February 1989.

[6] W3C, Data Integrity BBS Cryptosuites v1.0 — Candidate Recommendation Draft, 10 September 2026; not a final Recommendation.

[7] NIST, SP 800-63C, Digital Identity Guidelines: Federation and Assertions — Final Revision 4, 2025; §5.4 on subscriber-controlled wallets.

[8] OpenID Foundation, OpenID for Verifiable Presentations 1.0 — Final specification, 2025.

[9] FATF, Guidance on Digital ID — 6 March 2020.

[10] European Commission, Implementing Regulation (EU) 2024/2982 — 28 November 2024; EUR-Lex text.

[11] NIST, "Zero-Knowledge Proof," Cryptographic Glossary — Official terminology reference accessed November 2026.