

CHAPTERS 1-3

The Payment Parallel

Lessons from the Payments Industry That
Will Shape the Future of Identity Online

David Stearns






Table of Contents

About the Author	3
A Letter to Our Readers	4
Lessons from Payments: What Can the Identity Industry Learn?	6
Issuing Credentials: The Foundation of Payment Networks	11
Authenticating and Authorizing Transactions: Securing Identity Like Payments	21
About Proof	32





About the Author

[David Stearns](#) is a software engineer and historian of technology. He is the author of [*Electronic Value Exchange: Origins of the VISA Electronic Payment System*](#) (Springer 2011), which chronicles the evolution of VISA from a paper-based regional credit card program to a fully-digitized worldwide payment system. He has built software systems at Stripe, Adobe, Microsoft, and a few startups. In between software jobs, he earned his PhD in the Sociology and History of Technology from the University of Edinburgh (2008), and was a Senior Lecturer at the University of Washington for five years. He now writes on payments, identity, and software engineering.



A Letter to Our Readers

Dear Readers,

Proof, the world's first identity authorization network, is thrilled to introduce *The Payment Parallel*, a new series written by David Stearns, software engineer, historian, and author of the widely cited book on the history of VISA and payment networks: *Electronic Value Exchange: Origins of the VISA Electronic Payment System*.

It's no secret that trust is rapidly eroding online. Advancements in artificial intelligence have blurred the lines between real and fake, making it increasingly difficult to trust the people and things we rely on every day.

Much like the payments industry, which has undergone significant transformations over several decades, how we handle identity online is at a crossroads. The lessons we can draw from the history of payments provide invaluable insights that inform how to address our current challenges and shape the future of establishing trust on the internet.

We must begin to think of identity verification not as a one-off technical or regulatory necessity but as the cornerstone of establishing trust and security online. By understanding the parallels between the payments industry's evolution and the current state of identity solutions, we can uncover the strategies and innovations needed to overcome the siloed and fragmented nature of today's solutions.

Throughout this series, David will explore how the payments industry transformed from isolated, single-purpose solutions to a more integrated, interoperable, global network. He will examine the key milestones, technological advancements, and policy changes that paved the way for a more efficient, safe, and trusted payment ecosystem.

Proof believes that the future of identity must be network-based, and to drive that evolution forward, there is much to learn from how payment networks are designed and run today.

We invite you to join us in reading *The Payment Parallel* and discover how the lessons from the past can guide us toward a more secure and trustworthy digital future.

Sincerely,





Lessons from Payment

What Can the Identity Industry Learn?

I recently changed jobs, so I wanted to rollover my 401k from my previous employer into an IRA. I had previously set up an authenticated account with the plan administrator when I started that job, which I could use to adjust my contributions or change my investment choices, so I figured it should be relatively easy to initiate the rollover process as well. After signing in, I discovered that the process was a *little* more complicated than I expected, and surprisingly manual. I had to:

- 1 Download and print the forms on paper (thankfully I still had a printer—many people don't!)
- 2 Make an appointment with a notary public
- 3 Show the notary my state-issued ID so the notary could verify my identity
- 4 Sign the forms in front of the notary, who then manually stamped them and recorded the transaction in their journal
- 5 Mail the signed forms to the plan administrator through the postal service and wait

You might be thinking, “yeah, that seems reasonable considering you’re asking the administrator to empty your retirement account!” No doubt, such an action should require stringent security, but how secure is such a manual paper-based process?

- The notary is expected to spot fake or manipulated state-issued IDs with sophistication, and match a tiny outdated picture to the live person, even though they probably don’t have access to specialized verification equipment.
- The notary has very limited context with which to evaluate the probability of fraud—they have no access to the signer’s previous interactions with other notaries or the relying party.
- The plan administrator has no real way to validate that the notary stamp is authentic, or that it hasn’t been manipulated post-signing. They also lack a consistent way to verify that the notary behind that stamp is actually still licensed by their State and in good standing.
- The mailed forms could be intercepted and manipulated in-transit, post-notarization, with no real way to detect that.

It’s also striking that this process is still so manual and paper-based. It’s largely the same as it has been since 401ks were created in the late 1970s!

It’s also striking that this process is still so manual and paper-based. It’s largely the same as it has been since 401ks were created in the late 1970s!

Interestingly, this whole process is reminiscent of how payments were processed before electronic payment cards. Consider the process of purchasing goods with a paper check in the 1960s:

- 1 The consumer would fill out a paper form, instructing their bank to transfer funds to the merchant's account, and sign it in the presence of the cashier.
- 2 Since no online authorization system existed yet, merchants would typically require the consumer to produce a state-issued ID as well, unless the consumer was a trusted regular.
- 3 The cashier was expected to spot fake or manipulated state-issued IDs with sophistication, and match a tiny outdated picture to the live person, without access to any specialized verification technologies.
- 4 If everything looked legit, the cashier would see if the ID number was on a list of people who had written bad checks at that store before. If not, the cashier would write the number on the check so that it could be added to that list if the check bounced. These lists were not typically shared between merchants, so merchants couldn't know if that ID had already written bad checks elsewhere.
- 5 The merchant later deposited the check at their local bank, but there were no real protections against someone altering the check details before it got to the bank's check processing department.
- 6 The bank could compare the signature on the check to the signature used to open the account, but this was hardly a precise science, and a rather weak form of authentication.

But over the last 50 years the security and efficiency of payments have radically improved:

- Consumers now carry a standardized credential (a physical or virtual card), issued by their bank, that identifies them to the transactional network.

- Merchants can insert/tap the credential at the point of sale, positively authorize the transaction, and get a guarantee of payment within a few seconds.
- Since the network sees all transactions, and knows all past disputes, it can estimate the probability of fraud, and quickly cut off compromised credentials.
- Many networks can also authenticate the consumer via something they know (PIN or password) or something they are (biometrics).
- Chip cards and mobile device wallets (such as Apple Pay) also digitally sign the transaction details to prohibit tampering in-flight.
- When fraud does still occur, cardholders are automatically protected against financial loss, and the network provides a clear dispute resolution process.

This is not to say that payments are now perfect—to the contrary, they still have a long way to go—but it's striking just how much payments have improved over the last few decades, while processes like rolling over a 401k have largely stayed the same. Services such as Proof's Notarize make the notarization part easier and more secure, but a tighter and more complete digital integration would bring this closer to the convenience and safety of payments.


I think the identity industry could learn a few things from the history of payments. Many of our transactions that require legal identity verification are still similar to signing paper checks in the 1960s, but the same techniques developed by the payments industry could make those more efficient and safe, while simultaneously reducing friction for consumers.

Over the next few chapters we will learn a bit about how payments work and how they've changed over the last 50 years. We'll also muse about what it

would look like to apply various principles from the payments industry to transactional identity verification. Stay tuned!

Proof's Key Takeaways

- Payments have evolved dramatically over the past 50 years, but identity verification for major transactions remains largely unchanged
- Applying the same principles would create a network with:
 - A reusable, ubiquitous credential
 - Verifiable transactions
 - A network-driven approach to fraud prevention
 - A liability framework with consumer protections



Issuing Credentials

The Foundation of Payment Networks

Today I can walk into almost any merchant, in any part of the developed world, and pay for goods or services using nothing but a small piece of plastic. The merchant and I might have bank accounts in totally different banking systems, denominated in different currencies, and we might not even share a common language, but the merchant will happily accept my card as payment and let me walk out the door with valuable goods. All I have to do is tap or insert my card into the merchant's reader, which responds in a few seconds with an approval that *guarantees* my payment.

That card is essentially an economic identification. It contains a *credential*, issued to me by the bank that extends me credit. The form of that credential has changed a bit over the years—from numbers embossed on the front, to data encoded in a magnetic stripe on the back, to a cryptographic certificate embedded in a chip—but it's very similar to an identity credential. It identifies my account during a transaction, but it also vouches for my membership in a payment network that is trusted by billions of people worldwide. And it's that *network* the merchant trusts—not me and not even my issuing bank.

But it wasn't always this way. Just seventy years ago, most consumers used cash or paper checks, but the latter were effectively useless outside their local area because they weren't guaranteed. How did we get to today's world of general-purpose cards with embedded identity credentials and trusted electronic authorizations? And what might the identity industry learn from all of this? To answer these questions, we need to do a bit of history.

[A bank] card is essentially an economic identification. It contains a credential, issued to me by the bank that extends me credit.

Single-Purpose Cards

The First Step Toward a Payment Network

Americans have always purchased goods and services on credit. In the early days of the Republic there wasn't enough currency in circulation, and receiving money rarely lined up with when you needed to spend it. So most merchants kept credit accounts for their well-known customers, who would periodically settle up when they had access to enough currency.

By the early 1900s, the United States was rapidly urbanizing and merchants were expanding to multiple locations, so merchants could no longer rely on "analog facial recognition" to identify their credit account holders. Instead they gave their account holders small cards to carry and present to any cashier at any location. The card was a credential that identified the account, and was often signed by the account holder at the time of issue. This provided a rudimentary authentication mechanism, as the signature on the card could be compared to one made by the purchaser on the sales draft at transaction time.

As gasoline stations and other franchised business models spread across larger regions, they too issued cards to their regular customers, but now those cards could be used at any franchise location. Billing and collections was handled centrally by the franchiser, but most transactions were still unauthorized, as the amounts one could spend on gasoline in those days was small enough to justify the risk.

Narrow-Purpose Cards

Expanding Usability Across Merchants

Although these early cards could be used at multiple locations, they were still restricted to just one business or franchise. You could use your gasoline card at any service station with the same brand, but you couldn't use it to buy lunch on the road, or pay for a hotel room at your destination. This naturally limited the transaction volume these card programs could achieve.

This started to change in 1950 with the introduction of The Diner's Club, which was a single card that could be used at a wide range of restaurants and other businesses loosely-connected with the activities of "travel and entertainment." It was a credential issued primarily to businessmen (it was the 1950s, so most of them were indeed men) by the Diner's Club corporation, which also signed up the merchants to accept it. Merchants paid Diners around six to seven percent of every transaction in exchange for the increased business and guaranteed payment, and Diners collected the full amount from the cardholder, booking the difference as revenue.

Because it could be used at a much wider range of establishments, Diners grew much larger than the single-purpose cards that had come before it. At their peak, Diners had 1.3 million cardholders nationwide. But that was still

just the tip of the iceberg—their cardholder base was a small percentage of the roughly 100 million American adults at that time, and everyday consumer spending was orders of magnitude larger than business travel and entertainment. If an organization could offer a truly general-purpose card that could be used *anywhere*, its merchant network and transaction volume would quickly eclipse Diners and all the single-purpose cards combined.

General-Purpose Cards

The Breakthrough That Scaled Payments

Diners could have tried to expand into a more generalized payment network, but they remained focused on the domain of travel and business entertainment. They had also spent years advertising the card as a prestigious status symbol for the business community, so it would have been difficult for them to convince consumers it could be carried by everyone and used for everyday purchases (it wouldn't have been much of a "Club" if they let everyone in). Financing such everyday use would also require serious reserves, which a non-bank company like Diners simply didn't have.

But there was a bank that had the reach and resources to create such a card: the Bank of America (BoFA). It operated in California, which was one of the few states at the time that allowed banks to operate statewide. California was also one of the most populous and wealthiest states in the 1950s, and BoFA had a relationship with a majority of those residents, so it had access to enormous amounts of money. Despite being a very large commercial bank, it was also culturally disposed to get into the consumer credit card business, as it was founded by the son of an Italian immigrant who prioritized serving the everyday worker.

Chickens and Eggs

Solving the Multi-Sided Market Dilemma

By 1958 the BofA was ready to launch their BankAmericard, but they faced a common dilemma: how do they get consumers to apply for and use the card before merchants are willing to accept it, and how do they get merchants to accept the card before consumers are asking to use it?

This is a problem that all multi-sided markets face. Both sides of the market must be enrolled in the system at roughly the same time in order to bootstrap the network. Once enough of both sides join, it starts to produce a chain reaction that becomes self-sustaining, but getting the reaction started requires some kind of initial incentive or leverage.

BofA solved this problem by simply mass-issuing cards to most of their depositors without asking them if they actually wanted it. BofA also used their extensive branch network to sign up merchants, many of whom got their initial business loans from that branch, so they were inclined to keep their banker happy. BofA started with a controlled experiment in Fresno, but quickly expanded to all of California when they heard their rival banks were planning to do something similar.

In just one year, BofA issued 2 million cards and signed up 20,000 merchants, just in California. Amazingly, this was almost twice as large as the entire *nationwide* network built by Diners over the past decade.

Both sides of the market must be enrolled in the system at roughly the same time in order to bootstrap the network. Once enough of both sides join, it starts to produce a chain reaction that becomes self-sustaining.

But this unsolicited mass-issuance naturally came at a cost: fraud and defaults were rampant. Because processing at this time was entirely manual and paper-based, it took another few decades to really bring it under control (we will dive into this topic in much more detail in the next chapter).

The BankAmericard was ultimately successful, however, and was eventually licensed to other banks across the country. In 1970 the licensees formed an independent organization, of which BofA became a member, which later adopted the name VISA and expanded worldwide (for a much more in-depth history of VISA and their systems, see my book [Electronic Value Exchange](#)).

How Visa Created a Global Standard

VISA managed to put an identity credential into the pockets of most consumers in the developed world, so the identity industry could likely learn a thing or two from its history. I want to highlight two in this chapter, and we will discuss others in the subsequent ones.

First, identity verification, like payments, is a network that brings together consumers and relying parties for the purpose of completing transactions. These kinds of networks tend to grow super-linearly because the value of the network increases with the number of possible interactions between the participants. As more consumers and more relying parties enter the network, there are more possible opportunities to benefit from the network, making it even more valuable.

Because value growth is super-linear, networks that enroll more participants will gain an outsized advantage over those that remain smaller, and the number of participants will naturally be gated by the variety of use-cases the network can support. Just as in the case of payment cards, a

general-purpose credential will attract more participants and generate more transactional volume than a narrow-purpose one. That increased transactional volume will not only fuel the network's growth, it will also help it fight fraud (more on this in the next chapter).

But it's important to note that participants need to *think* of the network as general-purpose just as much as it needs to have the technical capability of being so. Most consumers in 1950s America thought of Diner's Club as an exclusive travel and entertainment card, not something for everyday purchases. The BankAmericard was advertised from the start as a general-purpose card, and VISA continued this with their "Everywhere you want to be" campaign.

To be more concrete, current digital identity credentials tend to be single-purpose or very narrow-purpose. They can get you through airport security faster, or authenticate you with a very small set of highly-related parties, but they can't be used to authorize a diverse set of transactions with a wide array of organizations. In this way they are like the store-specific or Travel and Entertainment cards of the 1950s—useful, but with limited reach. A truly general-purpose identity credential, which is used to authorize transactions with different relying parties than the one through which the person was initially enrolled, would be much more valuable, have potentially unlimited reach, and would quickly eclipse existing networks.

The second lesson this very abbreviated history shows is that multi-sided networks are tricky to start. They require strong incentives or some kind of leverage to enroll a critical mass of participants from both sides. The unsolicited mass-issuance approach ultimately worked for the

Identity verification, like payments, is a network that brings together consumers and relying parties for the purpose of completing transactions.

BankAmericard, but it also created a lot of fraud, primarily because their processing was all manual and paper-based at the time. If it had been electronic and online from the start they might have been able to keep that more under control.

But the Automated Clearing House (ACH), which is used for payroll direct deposit and electronic bill pay, provides a different and perhaps safer example. When it was launched in the early 1970s they convinced one of the largest payees in the nation, the United States Government, to use it for social security benefits and payroll. Consumers were told through TV commercials that they could get their benefits much more quickly and safely if they just opened a bank account and enrolled. Federal workers were encouraged to get their paychecks automatically deposited, saving them a trip to the bank.

But the ACH didn't stop with just the Federal Government—that would have kept them too narrow-purpose. Processing payroll efficiently and safely was a problem that all large organizations had, so they were able to use the Federal Government's successful rollout as motivation for most of the Fortune 500 companies to adopt the system as well. This not only raised public awareness of the system, but also encouraged smaller organizations to follow suit. Nobody was forced to enroll by mandate. The benefits were obvious, and the success of larger organizations convinced smaller ones it could work for them as well.

Building Identity Networks with Payment Strategies

VISA didn't succeed because they designed and issued the perfect verifiable credential, or the most secure wallet application. Instead, those early plastic

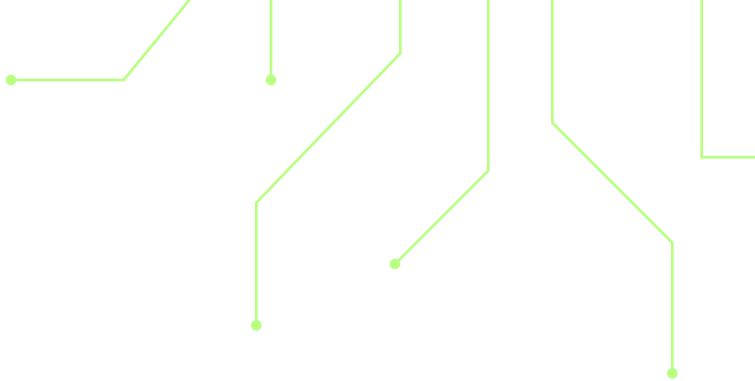
cards were just an imperfect mechanism for identifying consumers to their real goal: a global payments authorization network.

This network connects consumers with merchants to facilitate *guaranteed* payments, anytime and anywhere. It's this network that allows me to pay for goods or services at any merchant in the developed world. The specific credential I use to access this network is actually less important than the *authorization* this network and my credential issuer perform for each transaction. It's this authorization that guarantees the payment, and controls fraud. Without the network and the authorizations it facilitates, my payment card would be nothing but an untrusted piece of plastic.

In the next chapter I'll dive into the details of card payment authorization and how VISA was able to manage fraud as they rapidly expanded their network. We'll see how VISA's experience offers some valuable lessons for how we can build identity authorization networks.

Proof's Key Takeaways

- Everyone already carries an economic credential: a payment card
 - Early on, banks faced a chicken-and-egg problem—consumers wouldn't adopt cards without merchant acceptance, and merchants wouldn't accept them without consumer demand
 - BankAmericard tackled this by mass-issuing cards to consumers
 - Over time, cards evolved from single-use (T&E) to general-purpose, driving network effects and rapid growth
 - The technology improved—cardboard to plastic, magstripes to chips, and now biometric wallets—but the purpose remained the same
- What can we learn?
 - General-purpose identity credentials face the same adoption challenge. Mass issuance could work but may not be the best strategy
 - A better model is ACH, which grew through adoption by a large payor (e.g., the U.S. Government)
 - Generalized identity credentials will scale faster than specialized ones, creating stronger network effects, richer data, and better fraud detection
 - Identity credentials can leapfrog analog systems and go straight to advanced technology



Authenticating and Authorizing Transactions

Securing Identity Like Payments

VISA is the largest payment network in the world. In 2024 they processed over [15 trillion dollars](#) in payments, from 4.5 billion cards in over 200 different countries. Despite this rather colossal volume, VISA's losses due to fraud are probably lower than you might think: [less than one tenth of one percent](#) of their overall volume.

This isn't to say that payment networks like VISA are perfect—far from it—but they have managed to scale up their volume quite dramatically while keeping fraud at manageable levels. That means there are lessons to be learned. In this chapter we will examine how payment networks like VISA manage fraud and how similar techniques could be used in transactional identity verification.

How Payment Authorization Evolved

Although VISA's fraud rate is manageable today, it wasn't always that way. In the early days fraud rates were much higher—so much so that

most observers thought the card program would likely fail. But a series of innovations in the 1970s and 80s enabled VISA to stem fraud losses and create efficiencies that would allow the network to scale rapidly during the subsequent decades.

The card network we know as VISA today started as a privately-issued card of the Bank of America (BoFA), initially launched in 1958. This was during the heyday of Travel and Entertainment cards, such as Diners Club (started in 1950) and American Express (also launched in 1958), but BoFA wanted to offer consumers something more general-purpose and flexible: the BankAmericard could be used for everyday purchases, at many different kinds of merchants, and consumers could optionally choose to finance their purchases over time.

But consumers weren't really demanding such a payment device, nor were merchants particularly eager to pay the fees associated with accepting the card, so the BoFA knew they would have to do something dramatic to kick-start the program. As we saw in the previous chapter, their solution was to mass-issue unsolicited cards to all of their account holders in good standing, which was about half the adult population of California at that time.

Early Sources of Fraud

Lessons from the First Credit Cards

This approach worked, but it also created *a lot* of fraud. Consumers didn't even know the cards were coming, so mailbox thieves and unscrupulous postal carriers could take the easy-to-identify envelopes and use the cards for at least a month before the consumer or the bank even knew what was happening. Organized crime also stole blank card stock to create their own

counterfeit cards, and fraudulent merchants submitted bogus transactions. Within the first 15 months, BofA reported that the system had already lost \$8.8 million, but industry observers suspected the losses were closer to \$20 million (equivalent to about \$217 million today).

It was difficult for the BofA to stem this fraud because the system at this time was almost entirely manual and paper-based. Merchants were required to authorize the transaction only if the amount was above a particular threshold for that type of merchant, which was known as the “floor limit.” If the amount was above the floor limit, the merchant had to telephone a local authorization center, read the card number and transaction details over the phone, and wait while the authorizer flipped through paper reports to determine if the transaction should be authorized. This could take several minutes, so merchants often skipped this step if the customer looked reputable, or especially annoyed.

To complete the transaction, merchants manually filled out paper sales drafts with the transaction details. These drafts consisted of several layers of thin paper with sheets of carbon in-between, which transferred the information (albeit faintly) to the lower layers. The customer got one layer, the merchant retained another, and then deposited the last layer with the bank. This last layer was actually an IBM 80-column punch card, but the bank staff had to manually punch the transaction details into the card before it could be sorted and tabulated for the cardholder’s statement. Customers received their statement once a month, so they wouldn’t notice a problem until at least a month after the first fraudulent transaction, and perhaps longer if draft processing fell behind.

As customers discovered and reported fraudulent transactions, BofA started to publish lists of “hot cards” that merchants were asked to check at the start of a transaction. Since these lists grew quite long, checking not only added

another delay but also risked offending the cardholder, so many merchants would skip this unless the cardholder looked especially suspicious. Even if they did check, new lists were published on a fixed schedule and had to be mailed to merchants, so it would still take several days before merchants could learn about a newly compromised card.

Franchising and Interchange

Connecting a Growing Network

By the early 1960s BofA had weeded out enough stolen cards and bad credit risks to start making a profit despite continued fraud, so they

decided to franchise the program to other non-competing banks in other states starting in 1966. These franchisee banks experienced similar fraud issues while launching their programs, but the licensing network created some new problems for authorization and clearing. If a business traveler from California used their card at a hotel in New York (which banked with a licensee bank and not the BofA), how would the hotel obtain authorization, and how would the sales draft get routed back to the correct issuing bank?

By the early 1960s BofA had weeded out enough stolen cards and bad credit risks to start making a profit despite continued fraud.

Since the system was still manual and paper-based, the solutions were very inefficient and cumbersome. Interchange authorizations required what was known as a “two-legged call.” The merchant called their local authorization center, which then had to turn around and call the issuer’s authorization center and relay the information verbally. Interchange authorizations could

take several minutes, causing many merchants to just skip authorization if the customer didn't look terribly suspicious.

Routing the sales drafts back to the issuing bank created further delays in not only clearing but also presenting the transaction on the cardholder's statement. A fraudulent interchange transaction could take months floating through the system before it was discovered by the legitimate cardholder.

Online Authorization

The Game-Changer for Payment Security

By 1968 the BankAmericard system was in such a state of disarray that the franchisees revolted and organized their own independent organization to run and improve the system, convincing the BofA to relinquish the brand and join as member with only a bit more power than the rest. This organization was initially known as National BankAmericard, Inc, but was later rebranded as VISA in 1976.

The leader of this new organization, Dee Hock, realized that the system wouldn't be able to keep fraud under control as it grew unless the authorization process was entirely automated and required for all transactions. In 1973 he debuted an online authorization system that all local authorization centers across the country could use to get a (more or less) instant response, even for interchange authorizations. But merchants still had to read transaction details over telephone, so in the late 1970s and early 80s, VISA made the cards machine-readable via a magnetic stripe, and piloted cheap countertop devices that could transmit the stripe data to the authorization center electronically. A lower processing fee encouraged merchants to adopt these devices, and as they did, both the floor limits and

hot lists disappeared. All transactions were now positively authorized by the card issuer within a few seconds.

This led to a dramatic reduction in fraud—merchants participating in the initial pilot saw an average 82% decrease. Issuers could now block a compromised card as soon as it was reported lost or stolen, and they could detect obvious fraud patterns that had been difficult or impossible to see before. Over time issuers were able to build up statistical profiles of “normal use” and flag transactions for review that seemed to deviate from that profile. Since VISA’s online authorization system sees every transaction for every issuer, VISA itself was also able to develop its own fraud scoring algorithms, which issuers can use as yet another signal.

In retrospect, all of this looks painfully obvious, but at the time, it took some real determination. The basic ideas had been floating around for a couple of decades, but nobody had been able to make it work at any significant scale. Most of the requisite technology existed in the early 1970s, but no one had put it together to create a *nationwide* online authorization network. Several competing efforts were happening around the same time, but Visa was the first one to actually put it into production across the entire country.

Online authorizations, as well as electronic clearing through a centralized clearinghouse (also implemented in the 1970s), allowed VISA to control fraud as they grew, but they also had to continuously stay ahead of the fraudsters by improving the security of their credentials. Magnetic stripes made the cards machine-readable, but simple audio equipment could be used to skim the contents and create a counterfeit copy. Eventually all the major networks shifted to chip cards, which generate transaction-specific one-time use codes that are digitally signed using embedded certificates. In many regions, these chips also verify a PIN entered by the cardholder at the time of sale. To better secure online transactions (which can’t leverage the chip) the

networks also created protocols for password authentication (3DS). And as consumers adopted smart phones with biometric sensors, mobile wallets like Apple Pay further secure transactions using biometric authentication. Fighting fraud requires constant technological innovation, but that becomes easier when participants can cooperate via a network, and networks cooperate with each other.

Applying Payment Authorization Techniques to Identity

In the opening chapter of this series I described my recent experience rolling over a 401k, and I was surprised at how little the process has changed over the last 50 years, despite the rather significant improvements in payments over that same period. The process I went through resembled the BankAmericard system of the 1960s far more than a VISA payment in the 2020s.

Even though payments have shifted more online, authorizing something with your legal identity has remained mostly manual, offline, and paper-based. Information moves very slowly in offline paper-based operations, so when fraud occurs, it takes a long time before it is detected. Blocking credentials previously used for fraud becomes almost impossible without some kind of online network.

Even though payments have shifted more online, authorizing something with your legal identity has remained mostly manual, offline, and paper-based.

What if these authorizations worked more like modern payment networks? What if there was an online system that many different kinds of

organizations could use not just to authenticate the current user, but *actually* authorize the transaction? Such an identity authorization network would:

- 1 Issue digital credentials that are bound to verified legal identities. Like a payment card, these digital credentials would be usable with a wide array of relying parties, for various kinds of transactions that must be tied to a legal identity: signing documents, transferring assets, making claims, applying for services, etc.
- 2 Authorize those transactions using those previously-issued credentials. This goes beyond simply confirming that the current user is who they say they are—it involves handling the actual signing of agreements in a legal, compliant, and provable way. The issuer is literally guaranteeing the transaction’s integrity and enforceability.

Issuing a new credential will naturally need to be a high-friction process. Given the advancements in generative AI, it would likely require a human in the loop working with advanced technology designed for detecting these deepfakes. But that friction and operational expense will get amortized over all the subsequent uses of that credential. The more opportunities consumers have to use those credentials, the less onerous the process will seem, and the more valuable the credential will feel.

Because all subsequent transactions would be authorized, this sort of system would naturally see all the activity related to that credential, so it could better spot obviously fraudulent patterns, and immediately block signings involving credentials reported as compromised. As transaction volumes increase, more data would be available for training machine learning models that could flag suspicious

[This system] would naturally see all the activity related to that credential, so it could better spot obviously fraudulent patterns, and immediately block signings.

transactions, and route them into a higher-friction authorization experience (e.g., additional authentication, or even a human in the loop).

Authorizing all transactions through the network also allows the network or issuer to notify the credential holder about any activity involving that credential. For example, most banks will let you sign up for notifications about every authorization performed on your card, but the same doesn't yet exist for something like your driver's license, or the digital equivalent of that. If consumers were notified when their digital identity credential was used to take out a loan or sell property, they could quickly respond if it was a fraudulent usage.

This sort of system could also immediately notify relying parties when a credential is reported as compromised or misused, allowing them to potentially rollback or invalidate transactions that were actually fraudulent. And if a relying party discovers a fraudulent use after the fact, the system would make it easy for them to report that so the credential involved can be revoked, and other affected relying parties can be notified.

Lastly, it could become a competitive advantage if a system like this could offer consumers protections from fraud. When a VISA cardholder sees a change they didn't make, they can call their bank and contest it. The charge is then removed from their balance while the issuer follows a well-defined network process to resolve the dispute. But if someone uses a fake identity document to fraudulently cash out a retirement account or transfer ownership of an asset, the rightful owner has to fight it at their own expense, often through a lengthy court battle.

Whether there would be one central issuer and authorizer of these credentials for the network (like Discover or American Express) or multiple following a common set of standards and rules (like Visa or Mastercard) is

a topic for the identity industry to negotiate, but both models successfully co-exist in payments. The acquiring processors that serve merchants have connections to all the various card networks and can switch transactions to the appropriate one based on the credential used. The same could be done for identity authorizations.

Identity Authorization Reimagines Consumer Protection

I hinted above that an identity authorization network patterned on payment card networks would also protect consumers when fraud occurs. But what exactly does that look like? How does it work in payment card networks, and how might it be structured and funded in an identity authorization network? We will turn to this topic in the next chapter.

Proof's Key Takeaways

- Payment card use has surged while keeping fraud in check—but this required deliberate effort
 - Early authorizations were manual and only needed for transactions above the floor limit, making fraud easy and undetectable
 - In the 1970s and '80s, VISA digitized authorizations and clearing, reducing fraud while scaling rapidly
 - Chip cards and PINs secured in-person transactions, while 3DS added authentication online
 - Biometric wallets like Apple Pay further strengthened security across both channels
- What can we learn?
 - Transactional identity verification today resembles payments in the 1960s—paper-based, reliant on ink signatures, and prone to manual errors
 - It must be digitized, with transactions tied to a credential linked to legal identity



About Proof

Proof is the trusted platform for securing the digital economy. As critical commerce shifts online, Proof empowers companies to verify who is behind every digital interaction. From pioneering remote online notarization to launching a comprehensive identity authorization network, Proof secures transactions with industry-leading compliance and fraud prevention. Trusted by more than 7,000 organizations across financial services, government, real estate, and healthcare, Proof is defining trust in digital transactions. For more information, visit www.proof.com.

© 2025. Notarize, Inc. (dba Proof.com)