



RESEARCH · INFRASTRUCTURE

Offline authorization is not resilience — it's the foundation of it

European payment networks are mandating offline card capabilities across the continent, but the real challenge lies not in the terminal's ability to approve locally, but in how issuing infrastructure, issuer processors, and merchants manage the gap between the moment of authorization and the eventual settlement of accounts.

Franco Di Pietro · The Payments Corner — October 9, 2026 · 14 min read

Research

Always connected.
Always available?

The Payments Corner

thepaymentscorner.com

A familiar technology with a new purpose

I rarely write about an industry announcement on the day it comes out. Most of them read differently a week later, once the coverage settles and the details surface.

This one I wanted to write about right away.

Reading Kelly Devine's progress report on Mastercard's commitments to Europe's digital economy, I expected the familiar themes: localized processing, European data governance, and the second building at Mastercard's European Technology Hub in Dublin. All of it matters, and all of it builds on the roughly €250 million she committed a year ago to three new data centres in France.

The part I read twice was much shorter:

Offline payments.

It's an odd turn for our industry. For sixty years, since Mastercard's predecessor was established in 1966, we have invested extraordinary resources in making card payments faster, smarter, safer, and more connected. We moved from paper authorization to electronic networks, from magnetic stripes to chips, and from plastic to tokenized credentials on our phones. We built systems that can evaluate a transaction against enormous amounts of data and respond in milliseconds, almost anywhere in the world.

And now one of the most consequential developments in European payments is about making transactions work when those connections disappear.

Here's how I read it.

Haven't we been here before?

Stand-in authorization is not new. For decades, the networks have been able to approve or decline on an issuer's behalf when the issuer's host, or its processor, cannot respond. The principle is simple: the temporary failure of one participant shouldn't automatically stop commerce.

But stand-in has one non-negotiable requirement. The transaction must still reach the network.

When it can't, the fallback is also decades old. In the 1990s, Europay, Mastercard, and Visa developed what became the EMV specifications, and they built a distributed authorization model into them. The chip and terminal can authenticate the card

offline, verify the cardholder with an offline PIN, and run risk management locally. The card's own logic then decides whether to approve offline, decline, or insist on going online. Floor limits, consecutive-transaction counters, and cumulative amount limits keep that local authority bounded.

None of this was invented in 2026. Nordic and Baltic markets have been working on offline card payments in earnest, and Mastercard cites live implementations in Denmark, Sweden, Estonia, and Latvia.

So what's actually new?

The technology is thirty years old. The commitment to deploy it everywhere is not.

Under the requirements Mastercard announced on October 8, all newly issued Mastercard cards across Europe will be enabled for offline payments from February 2027, and new and replacement terminals must support them from May 2027. Eligible transactions can be approved locally, within limits the issuing bank sets in advance, up to €200 per transaction.

Acquirers should notice one discrepancy. Mastercard's own communications describe the terminal scope slightly differently: the press release refers to all new and replacement terminals, while Devine's piece points to new terminals at critical retailers such as supermarkets and petrol stations. Acquirers will want the rules text before planning around either reading. Either way, the transition will take three to five years, paced by card reissuance and terminal replacement cycles.

I find that encouraging. The industry is taking a capability it has had since the 1990s and treating it as baseline infrastructure rather than a contingency for rare cases.

Visa has been moving the same way.

When both global networks start emphasizing the same thing

In October 2025, Visa Europe's CEO, Antony Cahill, published a perspective on strengthening Europe's digital payment resilience. Beyond Visa's established stand-in

capabilities, he described working with issuers, acquirers, and regulators to enable offline payments during emergencies, backed by appropriate safeguards and liability arrangements.

In May 2026, Visa followed with a €500 million commitment to Europe over the next decade. It is anchored by a new Eurozone data centre and also includes a Frankfurt headquarters with a European Innovation Centre and a technology centre in Warsaw.

Mastercard's own commitments follow a similar line. A year ago, Devine committed roughly €250 million to three new data centres in France. Today's update describes localizing critical payment operations, keeping the primary copy of European transaction data in Europe, and aiming to process European transactions primarily within Europe by 2030.

These programs differ in scope, duration, and content, so comparing the headline amounts directly would be misleading. The direction is the same, though. Both networks are responding to the expectation that critical payment infrastructure must be more resilient, more locally accountable, and less dependent on any single point of failure.

For decades, we discussed resilience mostly in terms of infrastructure availability: redundant data centres, network routing, recovery-time objectives, service levels. Those remain essential.

But the network can be running perfectly while the merchant cannot reach it. A network can report five-nines availability while an entire community is unable to pay electronically.

Network availability and payment availability are not the same thing.

That gap is what the rest of this piece is about.

Three different ways to keep a payment moving

Before going further, it's worth separating three mechanisms that often get grouped under "offline payments." They respond to different failures, and they allocate risk differently.

	NETWORK STAND-IN	OFFLINE EMV AUTHORIZATION	STORE-AND-FORWARD
What has failed	The issuer's host or its processor	The terminal's connection to the network	The terminal's connection to the network
Who decides	The network, under stand-in parameters agreed with the issuer	The chip and terminal, under limits the issuer personalized onto the card	Nobody at the point of sale; the issuer decides later
Approved at the moment of purchase?	Yes, online	Yes, locally	No; the merchant accepts it on trust
Where decline risk usually sits	Defined by stand-in rules	Generally with the issuer, within its own parameters	Generally with the merchant

The middle column is the one this announcement concerns, and its distinction from the third matters most. A locally approved EMV transaction is not a merchant accepting a payment that may later be declined. The card made an authorization decision, under rules the issuer set in advance.

Even genuine offline EMV, however, is not a disconnected settlement system. The authorization happens locally, but the transaction must still be cleared, settled, and posted once connectivity returns.

And posting is someone else's job.

The issuer processor: the participant we rarely talk about

Network announcements rarely mention this part of the chain.

Behind every card sits an account, and behind that account sits issuing infrastructure that maintains balances, records transactions, enforces controls, and makes authorization decisions. Depending on the issuer's architecture, those functions live with an issuer processor, the bank's own core systems, or some combination of providers.

In credit, we call the remaining available credit the open-to-buy, or OTB. In debit, the equivalent is the available balance, adjusted for holds and authorization rules. An online authorization evaluates each purchase against that live position, together with recent authorizations, fraud signals, account status, and spending controls.

An offline approval happens without that live position. The card has to decide without knowing the customer's current open-to-buy.

The chip does not carry a synchronized copy of the account. What it carries is delegated spending authority, expressed as offline risk parameters. Depending on the card application, those include lower and upper consecutive offline limits, cumulative offline amount limits, per-transaction ceilings, and the counters that track usage against them. Within the lower limits, the card may approve offline. Past the upper limits, it must go online or decline.

The issuer processor can't authorize anything while the card is offline. What it controls is the setup, meaning the conditions under which a disconnected authorization can safely happen.

Before any outage, the issuer has to define its offline risk appetite. How much exposure is acceptable per card and across the portfolio? Which products are eligible? Should a long-tenured cardholder carry the same offline limits as a new account, or one already close to its credit line? Those answers have to be translated into card profiles, personalization data, authorization policy, and operating procedures, and the processing platform has to support each of them.

The practical difficulty is that a processor cannot refresh a chip whenever the account changes. Offline parameters are set at personalization and can be updated afterward only through supported post-issuance mechanisms, mainly issuer scripts delivered in the response to an online EMV transaction. In many card profiles, the offline counters are also reset only when the card completes an online transaction and validates the issuer's cryptographic response.

Both depend on the card still being in the terminal when that response arrives. Visa's Quick Chip approach in the United States lets the cardholder remove the card before the issuer responds, and it does not process issuer scripts. It speeds up checkout considerably. But in a Quick Chip flow, neither scripts nor issuer authentication reaches the chip.

So even with every network up, issuers face a question: how do issuing systems, card profiles, and terminal implementations keep a card's offline authority aligned with the account's current risk?

An actual offline transaction makes it harder still.

The €200 question: who really carries the risk?

Here is how it can go wrong.

A cardholder has €300 of open-to-buy. Her chip is personalized to allow qualifying offline transactions within the issuer's limits.

During a telecom outage in her town, she buys €100 of groceries. The chip and terminal approve it locally.

That evening, her household's connected spending continues: an online order and a recurring bill, €260 in total, both authorized normally against the live ledger. The issuer approves them because, as far as its systems know, €300 is still available.

When connectivity returns and the €100 offline transaction posts, her open-to-buy is €40. The account is €60 over limit, and there is nothing left to decline. Every

decision along the way was correct under the rules that applied to it. The exposure came from the gap between two correct decisions made on different information.

Across thousands or millions of cards, over a disruption that lasts days rather than hours, those gaps add up.

This is why Mastercard's €200 figure should not be read as a €200 offline allowance. It is a per-transaction ceiling. A card's real offline exposure is set by the issuer's cumulative and consecutive limits, and portfolio exposure is that figure multiplied by every card configured to use it.

Who absorbs the resulting loss depends on scheme rules, the circumstances of the authorization, and the agreements among participants. As a general principle, an offline approval made within parameters the issuer itself personalized onto the chip tends to be the issuer's risk to carry, which is why issuers need to set those parameters carefully. But there is no universal rule that every offline loss lands in one place, and nobody should plan as if there were.

Most of the cleanup falls to the processor. Once connectivity returns, the issuing platform must receive the offline records through clearing, post and reconcile them against the account, update the open-to-buy, and handle over-limit cases, exceptions, and disputes. Where the technology allows, it must also restore or adjust the card's offline controls.

An offline approval creates work the real-time authorization system couldn't do at the moment of purchase. That work doesn't disappear because the customer walked out with an approved receipt.

Offline payments distribute the authorization decision. They don't eliminate the need for centralized financial accountability.

For anyone designing issuing platforms, that is the main architectural lesson of this mandate.

Resilience requires more than enabling an EMV feature

Offline resilience can look like a card configuration change plus a terminal software update. In practice it depends on several parties, each owning a piece.

Issuers need explicit offline risk policies. Issuer processors need to support the account controls, card-management capabilities, and reconciliation those policies require. Personalization bureaus need to deploy the right chip applications and parameters at reissuance. Acquirers and merchants need compatible terminals, staff procedures, and clear liability terms. Networks provide the rules, interoperability, security requirements, and clearing that hold it together.

And not every card, terminal, or wallet behaves the same way. Contactless cards and mobile wallets follow their own risk logic, so a cardholder who taps a phone during an outage may not have the same experience as one who inserts a card and enters a PIN.

There is also the matter of power. Offline authorization does nothing for a merchant whose terminal has no electricity. The point-of-sale system, the payment device, and often the store's lighting and refrigeration all need power. Generators and batteries can bridge part of that gap, but only if someone planned for it before the emergency.

Mastercard's own example shows this. During a recent storm in Latvia, card payments continued in shops powered by local generators. The offline capability mattered. So did the generator.

Resilience is partly a software question and partly an operational one, involving the electricity provider, the telecom carrier, the merchant, the acquirer, the network, the issuer, and the processor. No single participant can solve it alone.

Europe is approaching resilience from another direction, too

The renewed focus on offline payments is unfolding alongside a larger European conversation about digital sovereignty: where transactions are processed, how data is governed, and who controls critical payment capabilities.

The ECB's digital euro is one expression of that ambition. Unlike offline card authorization, the digital euro is designed to support payments in prefunded digital central bank money, including device-to-device transactions without continuous connectivity.

The two work very differently. A locally approved EMV credit transaction is still a promise of payment that has to be cleared and settled later, against an account whose position may have moved. An offline digital euro payment aims to transfer value that was funded in advance. The customer experience might eventually look similar; the risk model underneath is fundamentally different.

The digital euro has not yet been issued. The ECB is preparing a pilot from the second half of 2027, with potential first issuance in 2029, subject to legislation and further decisions.

Then there is Wero, the European account-to-account initiative. Wero already offers a European alternative and is working toward broader point-of-sale acceptance. On October 8, the European Payments Initiative announced it had joined nexo standards to support interoperable acceptance. Wero currently requires connectivity, so it shouldn't be presented as an offline solution. But it reflects the same European interest in reducing structural dependencies and widening payment choice.

Mastercard, Visa, the digital euro, and Wero use different technologies, ownership structures, and financial models. They all operate in a market where resilience, control, interoperability, and infrastructure independence carry more weight every year.

I wouldn't read the networks' investments as simple reactions to the digital euro. Their resilience programs have their own history and business logic. But the competitive context is clearly changing.

To remain truly global, payment networks may increasingly need to prove how well they operate locally.

The same question looks different from the United States.

What about the United States?

Europe has an explicit sovereignty agenda. The United States doesn't approach payment infrastructure from the same geopolitical position. On operational resilience, though, the U.S. has plenty at stake.

Consider Hurricane Helene in September 2024. Across western North Carolina, damage to power and telecommunications infrastructure disrupted economic activity. Businesses that managed to reopen sometimes accepted only cash because their terminals couldn't connect. Many ATMs were offline or empty. Carolina Public Press described communities where residents had money in their accounts but struggled to reach it for essential purchases.

After a disaster, a family goes out to buy food, fuel, or medicine. The store is open, the shelves are stocked, and they have the money. The sale still fails, because the terminal can't reach the authorization network.

That isn't a problem of purchasing power.

It's a problem of access to purchasing power.

In an emergency, the second problem is the one people actually face.

The U.S. ecosystem has sophisticated network resilience: stand-in, redundancy, recovery procedures, and merchant-level deferred authorization. But American chip acceptance was built differently from Europe's, and the difference sits on the card as much as in the terminal.

When EMV arrived after the 2015 liability shift, the U.S. market optimized for online authorization. Most U.S. credit cards are personalized to go online for nearly every transaction, typically with no cardholder verification or a signature rather than a PIN, and offline PIN support is limited. The European model Mastercard is now scaling assumes chip-and-PIN: a verified cardholder, and a card carrying meaningful offline authority.

That makes the U.S. gap a card personalization and cardholder-verification problem, not just a terminal problem. Building offline payment continuity here would

mean revisiting card profiles at reissuance, issuer risk policy, processor support for offline parameters and reconciliation, and checkout flows like Quick Chip that were designed for a world where the network is always reachable. That is a multi-year effort paced by reissuance cycles, so the work has to start well before the next storm.

The Federal Reserve has studied it. In August 2024, it published research on the potential role of offline payments in improving the reliability and resilience of digital payment systems. The paper also weighed the risks: fraud, double spending, delayed processing, security, and merchant liability.

The problem is understood. What's missing is a decision to treat offline payment continuity as an explicit part of disaster preparedness.

Hurricanes are the obvious case, but not the only one. Wildfires, floods, severe winter storms, prolonged telecom outages, cyber incidents, and major infrastructure failures can all cut a community off from everyday commerce. The need to pay for essentials doesn't disappear when a merchant's connection does.

From technical resilience to community resilience

The April 2025 blackout in Spain and Portugal is another data point. The ECB later estimated that physical card spending in affected Spanish areas fell by roughly 41 to 42 percent during the disruption.

That drop shows how much everyday commerce depends on systems most consumers never think about until they stop. It also raises a policy question: should access to essential commerce during emergencies be an explicit objective of payment infrastructure policy?

No single technology would cover it. Offline EMV, network stand-in, redundant communications, cash availability, backup power, and merchant contingency procedures all contribute. The right mix will vary by market, merchant category, transaction type, and risk profile.

An offline grocery purchase during an outage is not the same risk decision as an offline purchase of a high-value discretionary item. Banks and processors can, and probably should, distinguish between the two, whether through tighter limits outside essential categories where the technology supports it, or through different parameters by product and customer segment.

That suggests a more precise way to plan. Instead of asking whether a card or terminal supports offline transactions, we should ask which transactions need to stay available, under what conditions, for how long, and at what level of acceptable exposure.

It may also mean measuring success differently. Not only by network uptime and recovery times, but by the share of essential transactions a community can still complete during a real disruption. That metric would connect technical architecture directly to economic resilience, and it would give issuers, acquirers, networks, and policymakers a common number to plan against.

The next frontier may not be more connectivity

For decades, we've measured progress in payments by acceptance, speed, security, fraud prevention, convenience, and ever more sophisticated authorization intelligence. We've become remarkably good at all of them. We can connect a cardholder to an issuer thousands of miles away, evaluate the transaction against enormous amounts of data, and respond before the card is back in the wallet.

That's an extraordinary achievement. Mastercard's announcement, Visa's European commitments, and the wider sovereignty debate raise a different question. How well does the ecosystem keep supporting commerce when the information and infrastructure it normally relies on are temporarily gone? And how much of that capability should be built in before the emergency, rather than improvised during it?

For issuing banks, that means deciding how much financial exposure can safely be delegated to the chip. For processors, it means supporting those decisions with card-management infrastructure, account-level controls, delayed-transaction handling, and reconciliation. For networks, it means interoperability and consistent

operating rules. For acquirers and merchants, it means acceptance environments, and power, that can operate in less-than-ideal conditions. And for policymakers, it means treating payment access as part of economic continuity planning.

Real-time authorization stays. The goal is systems that know what to do when the information they depend on isn't there.

I started writing because I was struck by the irony of revisiting offline authorization sixty years into the history of modern card networks. After working through it, I think the story is less about the past than about what card networks have to deliver next.

Consumers don't care whether a transaction was approved by their bank, by a network in stand-in, or by a chip following rules set months earlier. They care that the payment works, especially when they need it most.

Maybe the next frontier of digital payments isn't making transactions faster, smarter, or more connected.

Maybe it's making sure they still work when everything else doesn't.

After all, being always connected and being always available are not the same thing.

Disclosure. The Payments Corner's founder is employed at Euronet Worldwide, a global infrastructure and transaction processing company. Analysis is category-level; nothing here is investment advice.

© 2026 The Payments Corner · thepaymentscorner.com/research/offline-authorization-is-not-resilience-its-the-foundation-of-it