Automated Fare Collection System Overview

TransitFare • 2026-08-23

A transit director evaluating three vendors for an automated fare collection system usually gets three different sales pitches and no shared vocabulary to compare them against. This whitepaper is that shared vocabulary. It breaks an AFC system into the four layers every implementation shares, regardless of vendor. It then works through fare media and validation, and covers what a deployment involves for an agency running a handful of routes rather than a few hundred. Where a decision is more complicated than a sales deck admits, that gets said plainly here.

What Is an Automated Fare Collection System?

An automated fare collection system combines fare media, validation hardware, back office software, and reporting tools. Together they let riders pay a transit fare without a driver or clerk processing each transaction by hand. Every tap, scan, or swipe becomes a record the agency can price, audit, and report against national transit standards.

The category has existed long enough to have its own standards effort. The American Public Transportation Association’s Universal Transit Farecard Standard, published in parts as the Contactless Fare Media System Standard, defines how contactless fare media should behave. A card issued by one agency should be read correctly by hardware built by a different vendor. Most small agencies never open the standard itself. Many contactless validators are built against some version of it, named on a spec sheet or not.

This piece describes that architecture the same way regardless of which vendor eventually builds it. An agency comparing proposals needs one shared framework, not three incompatible ones from three sales teams.

The Four Layers of an AFC System

Every automated fare collection system, from any vendor, is built from the same four layers stacked on top of each other. A system is only as strong as its weakest layer, and vendors tend to talk most about whichever layer they built first.

Fare media is what the rider actually carries or presents: a card, a phone, a paper ticket, cash. This is the layer riders see, and the one vendors market hardest.

Validation is the device at the door or the platform that reads the fare media, checks it against a fare table, and accepts or rejects the trip in real time.

The back office is the software that defines what a valid fare actually is: the products, the prices, the capping rules, the blocklists, the reconciliation against settlement.

Reporting turns every validation into the numbers an agency needs. That means ridership by stop and route, revenue by payment type, and the fields the National Transit Database (NTD) asks for every year.

A system that is strong in fare media and weak in reporting will collect fares just fine. It can still leave a general manager unable to answer a board question about which route earns back its subsidy. The sections below work through each layer in order, then cover what a deployment involves and what tends to go wrong.

Fare Media, Compared

Six fare media types cover nearly every automated fare collection system running in North America today, cash included. None of them is obsolete, and none of them is sufficient alone. Most small agencies run at least three at once during any transition.

- Fare medium | How value is stored | What the rider needs | Typical lifespan
- Cash | Physical currency, counted at the farebox | Coins or bills | Not applicable
- Magnetic stripe | Value or a trip count encoded on the stripe | A printed card, reissued as it wears | Weeks to a few months of daily use
- Contactless smart card | Value or a linked account ID on an embedded chip, per ISO/IEC 14443 | A reusable card, tapped rather than swiped | Multiple years of daily taps

The remaining three trade an agency-issued credential for something the rider already owns:

- Fare medium | How value is stored | What the rider needs | Typical lifespan
- Barcode or QR | A signed, time-limited code, app-generated or kiosk-printed | A phone screen or a printed slip | A single trip to a season, expiring by design
- Mobile wallet | An account balance or ticket held in the phone’s wallet app | A smartphone with the wallet app installed | As long as the phone stays charged
- Open-loop EMV | The rider’s own bank card or device, no agency-issued credential, per the EMV contactless specifications | A contactless bank card or a payment-enabled phone | Whatever the rider’s own card’s expiry is

Contactless smart card and open-loop EMV validators are usually the same physical reader, since both operate at 13.56 MHz under the same ISO/IEC 14443 proximity standard. What differs is the logic behind the tap. A smart card looks up an account the agency issued and controls locally. Open-loop EMV instead routes the transaction to a card network for authorization. A single validator can be built to read all of these on one device, though which ones a given vendor’s validator actually accepts is worth confirming.

Cash still needs a physical collection point even after contactless arrives. Most agencies keep existing fareboxes installed alongside the new validator rather than removing them. Cash keeps working through the transition, instead of forcing every rider to switch on day one.

Across TransitFare’s own AFC deployments, contactless smart card produced 3.5 million taps in 2025, ahead of over 1 million barcode validations that same year. That split is a rough proxy for how the mix tends to run at small agencies: cards first, then phones and barcodes.

Across the industry, open-loop acceptance is growing from a smaller base.

How Validation Works

Validation is the moment an automated fare collection system actually does its job. It reads the fare media, decides whether the fare is good, logs the result, and lets the rider board. Everything before that moment is setup. Everything after it is bookkeeping.

Two validation models cover almost every small and mid-sized agency. Onboard validation puts the reader at the farebox or a dedicated unit near the front door. The driver gets some visibility into whether a rider tapped. Platform or off-board validation, common on rail and bus rapid transit, puts the reader at the stop itself instead. It relies on proof-of-payment enforcement rather than a driver watching every boarding. Most fixed-route agencies running 10 to 50 vehicles use onboard validation. It needs no separate platform infrastructure, and it keeps enforcement with the person already standing at the front of the bus.

The case a lot of AFC marketing skips is the one rural and small-city agencies actually live in: what the validator does when it cannot reach the network. A validator that only works online stops accepting fares the moment a bus crosses a dead zone on a rural loop. For many small agencies that is not an edge case. It is most of the route.

A validator built for offline operation keeps a local copy of the fare table and the current blocklist. It carries enough logic to make an accept-or-reject decision without calling the back office first. Every transaction is logged locally with a timestamp. It then syncs the whole batch the next time it has a signal, whether that is at the next stop with coverage or back at the garage overnight. The rider experience does not change. The back office simply sees the transactions later than it would on a connected route. Reconciliation has to account for that lag, rather than assume every trip lands in real time.

Account-Based Versus Card-Based Architecture

Underneath fare media and validation sits a decision that shapes almost everything else in an automated fare collection system. It is where the fare rules and the rider’s balance actually live. Two architectures answer that differently, and neither is a wrong answer, only a different set of tradeoffs. A longer treatment of this decision is in account-based ticketing for small agencies; the short version follows here.

- Card-based | Account-based
- Where value and rules live | On the card’s chip, read and rewritten at every tap | In the cloud; the card or phone is only an identifier
- What a fare policy change requires | Updating logic on every validator, sometimes every card | One change in one place, live once validators sync
- What happens when a card is lost | The value on it is usually gone with it | Balance and history move to a replacement card

- Offline behavior | Naturally resilient; the card carries its own truth | Needs a local cache and a sync strategy, as above
- Typical fit | Agencies wanting the simplest validator and no cloud dependency | Agencies wanting fast policy changes, capping, and account recovery

Account-based systems make fare capping, mid-year concession changes, and lost-card recovery straightforward, since none of them require touching a single physical card. They also make a validator that has been offline for an extended stretch a harder problem, because the account balance it checks against may be stale. Card-based systems flip that trade: the card is self-contained and forgiving of a bad connection, but a policy change means touching every card the agency has already issued.

The Back Office

The back office is the layer riders never see and boards ask about constantly, since it is where fare policy actually gets enforced. Four jobs sit inside it.

Fare products. A fare table is rarely one price. A typical small agency back office manages a base cash-equivalent fare, a day pass, a monthly or period pass, and at least one concession tier. Federal grant conditions add a floor under that last one. An agency receiving Urbanized Area Formula Program funding must give certain riders an off-peak fare no higher than 50% of the peak fare. Those riders are older adults, people with disabilities, and Medicare cardholders. That rule is documented in FTA’s half-fare guidance. The back office has to encode it correctly for every fare medium at once, not only the one the agency issues most.

Capping rules. Fare capping lets a rider paying per trip automatically stop paying once the day’s or period’s taps add up to the price of a pass. The rider never has to buy the pass upfront. It is a back office feature before it is anything else. The logic has to track a rolling total per rider across every validator in the system, and stop charging at the right threshold, in real time, at every door.

Blocklists. A lost, stolen, or charged-back card needs to stop working everywhere, not only at the validator where it was reported. The back office maintains that blocklist and pushes it to every validator, including offline ones, on their next sync.

Reconciliation. Every validation has to match a settlement record. For agency-issued fare media that is mostly bookkeeping. Open-loop EMV involves an outside party. Transactions typically settle to the agency’s account in 1 to 3 days, once the card network and the agency’s payment processor complete their side of it. The back office has to reconcile a validator log against a bank deposit landing on a different schedule than the tap itself.

Reporting and What a Board Actually Asks For

A back office that gets fare policy right, paired with a reporting layer that cannot answer a board’s questions, is still a failed deployment. The board is the audience deciding

whether the system gets renewed.

Three questions come up at nearly every board meeting where fare collection is on the agenda. Which routes and stops carry the ridership? What share of revenue comes from cash versus card versus other payment types? And is the agency compliant? Revenue mix matters because it drives everything from armored car costs to processor fees. Compliance means correctly reporting unlinked passenger trips to the National Transit Database. For larger agencies, it also means producing the fare equity analysis a fare change requires under FTA Circular 4702.1B, the Title VI circular.

A reporting layer built for this has to do two things a spreadsheet exported from the back office usually cannot. It has to hold data long enough to show a trend across fiscal years. And it has to format that data the way NTD and Title VI reviewers actually ask for it, not the way the validator happened to log it.

What Does an AFC Deployment Involve for a Small Agency?

A deployment is not a single event. It runs in phases, and skipping one to save time is the single most common reason a launch date slips.

A twelve-week timeline, typical for TransitFare's AFC deployments including hardware, breaks into five phases. Most vendors' projects share these phases in some form, even when the labels or the compression differ:

- Phase | What happens | Typical duration | Who needs to be in the room
- 1. Discovery and fare policy design | Fare products, concession rules, and capping thresholds get finalized on paper before hardware ships | 2 weeks | General manager, finance, the vendor's implementation lead
- 2. Hardware procurement and configuration | Validators are ordered, configured against the agreed fare table, and bench-tested | 3 weeks | IT or operations lead, vendor hardware team
- 3. Back office and fare table setup | Products, blocklists, and reporting fields are built and connected to the agency's NTD workflow | 3 weeks | Finance, the person who currently owns NTD submissions
- 4. Installation and field testing | Validators go into vehicles and get tested against real routes, including any known dead zones | 3 weeks | Maintenance staff, drivers on the test routes
- 5. Go-live and staff training | Drivers, dispatch, and finance are trained on the new workflow, with the old cash process running in parallel briefly | 1 week | Every driver, dispatch, finance

None of this needs a dedicated project manager or an IT department of twenty. It needs one person on the agency side, usually the general manager or an operations lead. That person needs the authority to finalize the fare table in phase one, so phases two and three are not built against a moving target. Deployments that run long usually trace back to that decision

arriving late, not to the hardware itself.

An agency adding only back office and reporting to a validator fleet it already owns can move faster. TransitFare's software-only AFC deployments typically take about 4 weeks, since there is no hardware installation phase to work through.

What Goes Wrong

An AFC project can hit every milestone on the plan and still fail the agency using it. Five failure modes account for most of the trouble, across vendors, TransitFare's own early deployments included.

The fare table gets finalized after hardware ships. Validators get configured, installed, and tested against a fare policy that finance changes two weeks before go-live. The whole install gets redone. This is the single most preventable failure on the list, and preventing it sits entirely with the agency. Lock the fare table in phase one, and treat any later change as a change order, not a quick fix.

The offline case gets discovered in production, not in testing. A validator never tested on the real rural loop looks fine in a parking-lot demo. It fails the first week it runs a real route. Field testing has to happen on the routes with the worst coverage, not the easiest one.

Reconciliation gets treated as a launch-week task instead of a standing job. Someone has to own matching validator logs against bank settlement every week, indefinitely. If that ownership is not assigned before go-live, it quietly does not happen. Nobody notices the gap until an audit does.

A fare change ships without an equity analysis. Skipping the Title VI review before a fare increase risks more than legal exposure. It can force the agency to unwind a change after riders have already adjusted to it, worse than the delay of doing the analysis first.

The back office gets scoped for today's fare media, not next year's. An agency might add mobile wallet or open-loop EMV a year after launch. If the back office was never built to handle a second settlement source, that agency ends up redoing work. A slightly wider initial scope would have avoided it.

FAQ

How Much Does an Automated Fare Collection System Cost?

Total cost depends on the number of validators, whether the back office is hosted or self-managed, and how much cash handling stays in place alongside the new fare media. No single figure applies across agencies of different sizes. Ask any vendor for a full lifecycle figure covering hardware, software, support, and payment processing, not just a validator's unit price. The unit price alone rarely reflects what the system costs to run.

What Is the Difference Between Open-Loop and Closed-Loop Fare Collection?

Closed-loop fare collection uses fare media the agency itself issues: a smart card, a barcode ticket, an app-based pass tied to an account the agency controls. Open-loop fare collection accepts the rider's own bank-issued contactless card or phone through the EMV network, with no agency enrollment step at all. Some agencies run both at once, closed-loop for regular riders and open-loop for visitors and occasional trips. Most small agencies today run closed-loop only.

Can a Small Transit Agency Afford Automated Fare Collection?

Affordability depends more on scope than on agency size. A phased rollout, starting with validators and back office on the highest-ridership routes, spreads the investment instead of requiring a full-fleet cutover on day one. Grant funding through standard federal formula programs commonly covers a meaningful share of AFC hardware for small and rural agencies. That shifts the question from whether an agency can pay outright to how the procurement gets structured.

Does Automated Fare Collection Eliminate Cash?

No, and a vendor promising that outright deserves some skepticism. Cash remains the only fare medium requiring no enrollment, no smartphone, and no bank account, which makes it the fare medium of last resort for unbanked riders and visitors. A well-designed AFC system reduces cash's share of transactions and the labor of counting it. Removing it entirely is a policy decision separate from the technology, one with real equity consequences attached.

How Is Fare Collection Data Used for National Transit Database Reporting?

Every validation feeds into unlinked passenger trip counts, one of the core ridership figures the National Transit Database requires from formula grant recipients each year. An automated system produces that count directly from validator logs instead of a manual boarding estimate, which is more accurate. That accuracy only holds if the reporting layer maps its fields to NTD's definitions correctly, verified during deployment rather than after the first annual submission gets rejected.

Using This Framework to Evaluate a Vendor

None of the four layers above require a specific vendor to exist, which is the point of laying them out this way. Any proposal on a director's desk should be describable in these terms. One that cannot be does not deserve the fare table yet.

TransitFare is a transit technology company founded in 2012 that builds Automated Fare Collection (AFC) for small and mid-sized transit agencies. It also builds Transit Alerts, Real-Time Tracking & CAD/AVL, and Automated Passenger Counting, all on the same TransitFare Cloud platform. The TF4 Edge Gateway is the onboard hardware tying them together for agencies that add more than one module. Its AFC deployments have processed 60 million+ fare

validations since 2012, with a strong production reliability record. Over 90% of its customers expand to additional modules.

Whatever vendor an agency chooses, three questions are worth asking. How do the four layers actually connect? What does the offline case really do? What happens to a fare policy change on day two of a five-year contract? Explore Automated Fare Collection to see how TransitFare answers them, or download the whitepaper edition of this overview to bring into a vendor evaluation.