



ACF IN ACTION

acfstandard.com

ACF en Action applique l'Agentic Commerce Framework® à des entreprises réelles, à partir de sources publiques. Chaque cas montre qui reste responsable quand l'IA décide.

ACF-CS-019 · Case Study #019 · Édition 2026

Agentic payment rails

How do you let an agent transact autonomously without giving it an unlimited spending power?

CARTE D'IDENTITÉ

SECTOR Payments	COUNTRY International	TECHNOLOGY Purchasing agent	TYPE Agentic AI	AUTONOMY bounded autonomous
CONFIDENCE High				

OUTILS ACF ILLUSTRÉS DANS CE CAS

● ACF-01	● ACF-02	● ACF-03	● ACF-06	● ACF-00	● ACF-05
----------	----------	----------	----------	----------	----------

● mobilisé · ● à prévoir · 6 outils du Toolkit (4 mobilisés, 2 à prévoir, sur 17).

1. Executive summary

In 2025, agentic commerce became concrete: AI agents buy on behalf of the user. To frame this, the major payment networks launched dedicated mechanisms: **Visa Intelligent Commerce**, **Mastercard Agent Pay**, and protocols such as **AP2** (Agent Payments Protocol, Google). Their object: to allow an agent to pay, but within strict bounds.

The decision delegated to the agent is one of the most binding there is: spending money on behalf of a principal. Read through the lens of ACF, this case is the **"reference"** counterpart to the Replit incident. Both deal with the same object, **autonomy bounded by caps and a brake**, by opposite paths. Where Replit let an agent carry out an irreversible action with no bound, the payment rails make the drift technically **impossible** beyond a threshold. Trust does not rest on "the agent will not get it wrong", but on design bounds. This is the positive demonstration of what ACF calls bounded autonomy.

2. Context

When an agent buys for a user, it commits their money: it is a high-stakes decision, often irreversible. The intent of the payment networks is to make this delegation possible **in confidence**, without giving the agent an ordinary means of payment that would authorize it to spend without limit. The question of the case: how can an agent transact fast and alone while remaining unable to exceed what its mandate authorizes?

The answer rests on one principle: neither prohibit autonomy, nor monitor it passively, but set caps and reserve the binding for the human.

3. Public sources used

- Visa, "Visa Intelligent Commerce" (2025 announcements).
- Mastercard, "Agent Pay" (2025 announcements).
- Google, AP2 protocol (Agent Payments Protocol); press pickups (Electronic Payments International, Digital Commerce 360, 2025).

4. Reconstructed AI architecture (from public sources)

Agent-scoped payment tokens (*agentic tokens*), deployed by the payment networks from 2025.

- Nature: payment credentials fitted with programmable controls, distinct from an ordinary means of payment.
- Function: allow an agent to pay within the limits of a mandate defined by the user.
- Actual scope: the agent acts fast and alone within its mandate; it is technically impossible for it to exceed the cap or step outside the authorized categories.

The programmable controls (Exhibit).

Control	Detail
Spending cap	Maximum amount the agent can commit
Categories / merchants	Restriction to authorized categories or vendors, or an allow-list
Bounded mandate	Object, amount, and duration defined by the user
Human confirmation	Required for large or irreversible purchases, beyond a threshold

Graded autonomy (Exhibit).

Purchase type	Criticality	Handling
Small routine repurchase	Low	Autonomous, within the cap
Significant purchase	High	Human confirmation
Irreversible purchase	High	Explicit mandate required

The guiding principle. The agentic token does not ask the agent to behave well: it **prevents** it from exceeding. The cap is a technical bound, not an instruction. Trust rests on bounds that make the drift impossible beyond a certain point.

5. ACF mapping

 **Tools used:** ACF-06 (caps and kill switch) · ACF-01 (graded autonomy) · ACF-12 (agent mandate) · ACF-02 (criticality).

In ACF language, the mechanism reads as follows:

ACF element	In the agentic payment rails
Named human (accountable)	The user (the principal) who grants a bounded mandate and answers for their purchases.
Agent	The AI buyer agent, acting on behalf of the user via a scoped token.
Platform	The payment networks (Visa, Mastercard) and protocols (AP2) providing the tokens and controls.
Delegated decision	Spend money, within the limits of the mandate: amount, categories, duration.
Autonomy level	Graded: autonomous on small repurchases within the cap, bounded (human confirmation) on large and irreversible ones.
Non-delegable zones	Large or irreversible purchases: reserved for an explicit human confirmation.
Rule / safeguard	Technical bounds: cap, authorized categories, bounded mandate, confirmation beyond a threshold. The drift is impossible, not merely discouraged.
Control third party	The payment networks, which enforce the controls at the token level.
Traceability	Each transaction is tied to the mandate and its bounds, therefore reconstructable and enforceable.

A schema in ACF graphic language accompanies this case (separate file).

6. Analysis

 **Tools used:** ACF-06 (bounds by design) · ACF-01 (graded autonomy) · ACF-12 (mandate) · ACF-02 (criticality).

What the mechanism does well (the reference design):

- **The bound makes the drift impossible, not merely prohibited.** The agent cannot exceed the cap or step outside the categories, technically. This is a safeguard by design, not an after-the-fact surveillance. It is the exact opposite of Replit, where the agent could carry out an irreversible action with no brake.
- **Autonomy is graded by criticality.** The agent is very autonomous on small repurchases (low criticality, within the cap) and bounded on large ones (human confirmation). The autonomy level depends on the stakes of the purchase, not on the agent's capabilities.
- **The mandate makes the delegation possible.** The controls are not a hindrance: they are what **allows** a decision as binding as spending to be entrusted. An agent can be very autonomous **within** its mandate; it is the mandate that makes this autonomy acceptable.
- **The irreversible stays with the human.** Non-cancellable purchases require an explicit mandate or a confirmation, whatever the agent's capability.

What ACF names (and generalizes):

- an implicit **constitution**: caps, categories, duration, confirmation threshold, set before any action;
- a de facto **kill switch**: beyond the threshold, the action is blocked and escalated to the human;
- a natural **register**: each transaction is tied to its mandate, therefore traceable and enforceable.

The lesson is not "banks are better managed than Replit". It is a principle: an agent that acts in a binding way is governed by technical bounds and a human confirmation on the irreversible.

7. ACF toolkit used

This grid indicates the tools that the ACF analysis uses (✓) to read this reference case, and those to anticipate (→
SOON). It does not mean the actors use ACF: it shows which instruments the framework would bring to formalize what the mechanism already does.

Tool	Status	Role in this case
ACF-06 · Caps and kill switch	✓	Set the bounds that make the drift impossible beyond a threshold
ACF-01 · Autonomy map	✓	Grade autonomy according to the criticality of the purchase
ACF-12 · Agent mandate	✓	Bound the object, the amount, and the duration of the delegation
ACF-02 · Criticality matrix	✓	Classify purchases (small repurchase / significant / irreversible)
ACF-08 · Decisions register	→ SOON	Formalize the trace of each transaction and its mandate
ACF-05 · Supervision	→ SOON	Link the tracking of transactions to a continuous monitoring

8. Governance questions (what an AI committee should ask itself)

1. Why does a technical bound protect better than an instruction, when the agent spends money?
2. Does a bounded mandate restrict the agent, or does it make its delegation possible in confidence?
3. Which decisions must remain under human confirmation, even for a very capable agent?
4. Are our caps imposed at the technical level (token, means of payment), or are they only a circumventable internal rule?
5. Is each transaction of the agent tied to a traceable and enforceable mandate?

9. General lessons

An agent that spends money is not governed by trusting it, but by **bounding it by design**. The agentic payment rails show that you can give an agent strong autonomy without giving it an unlimited power: the cap, the categories, and the mandate make the drift impossible beyond a threshold, and the irreversible stays with the human.

The transferable lesson: bounding an agent by a technical constraint is not hindering it, it is making its delegation possible in confidence. You delegate execution within a mandate, never an unlimited power.

10. What ACF would do (reference design)

Prescriptive formulation: here is how an organization would design such an agent directly with the framework, whatever its sector.

If an organization wished to deploy this type of agent directly with ACF, it could in particular:

- **define** a bounded mandate (ACF-12): object, maximum amount, authorized categories, duration;
- **set** the bounds as a safeguard by design (ACF-06): cap, allow-list, human confirmation beyond a threshold, so that the excess is technically impossible;
- **grade** autonomy according to the criticality of the purchase (ACF-01 and ACF-02): autonomous on small repurchases, confirmation on large ones, explicit mandate on the irreversible;
- **document** a decisions register (ACF-08) linking each transaction to its mandate and its bounds;
- **link** this register to a continuous supervision (ACF-05) to detect any anomaly at scale.

This case is presented as a **reference design**: it shows, positively, what a well-governed binding agent must have, and serves as a contrast to the Replit case.

11. Similar cases (transferability)

This case is specific neither to payment nor to card networks. It is transferable to any organization that entrusts to an agent a **binding decision bounded by a mandate**:

-  agentic commerce and automated purchasing
-  finance and treasury (execution agents under a cap)
-  procurement and corporate purchasing
-  cloud services and infrastructure (agents that provision under a quota)
-  marketing and advertising (agents that commit a media budget)
-  any agent authorized to commit a limited resource on behalf of a principal

Everywhere, the same pattern: an agent that acts fast and alone **within** a mandate, a technical bound that makes the excess impossible, and a human confirmation reserved for the binding and the irreversible. This is what makes this document a **reference case**: it illustrates a concept, not just a mechanism.

12. Confidence of the assessment

High. The case rests on first-hand **official announcements** (Visa Intelligent Commerce, Mastercard Agent Pay, Google's AP2 protocol) and on specialized press pickups from 2025. The mechanisms described (scoped tokens, caps, categories, mandates, human confirmation) are therefore established and verifiable by the reader. The ACF analysis draws on these public facts and bears on the **governance logic**; the implementation details specific to each network, which are not public, have no bearing on the reading proposed here.

Teacher guide

Pedagogical objective. Convey that you bound a binding agent **by design**, and that the technical bound does not hinder autonomy: it makes it acceptable. Teach this case in contrast with Replit.

Session outline (90 min).

1. (15 min) Reading of the **two** cases of the kit: Replit (incident) and payment rails (reference), without the analyses.
2. (20 min) In groups: describe the bounds as a safeguard by design (card ACF-06).
3. (20 min) Exercise: define a bounded mandate and the graded autonomy of a buyer agent (card ACF-01).
4. (20 min) Comparison Replit / payment rails: instruction vs technical bound.
5. (15 min) Synthesis: you bound the agent by design, and you reserve the irreversible for the human.

Possible exam questions. See the "Governance questions" section.

Reminder: grading students' work is a pedagogical act of the teacher. This guide provides the reference analysis, not automatic grading.

Associated lab (ACF-LAB-019)

From this case, the student produces:

- the description of the bounds as a safeguard by design (card ACF-06): cap, categories, confirmation;
- the bounded mandate and the graded autonomy of a buyer agent (cards ACF-12 and ACF-01);
- the criticality matrix of the purchase types (card ACF-02): small repurchase, significant, irreversible;
- the one-page comparison with the Replit case: what changes on the side of caps and human confirmation.

Appendix - Completed ACF cards

The official Toolkit cards used in this case, filled in with its public data. For illustration: an organization of the same profile can reuse and adapt them.



OBJECTIVE

Map the purchases delegated to the agent and grade autonomy by criticality, not by the agent's capability. Filled from public sources (Visa, Mastercard, AP2).

1 AGENTIC MATURITY SCALE

0

Classic automation

Fixed rules, no autonomy.

1

Assisted agents

The agent proposes, the human decides.

2

Governed agents

The agent acts within a frame, under supervision.

TARGET ~80%

3

Supervised autonomy

The agent decides alone; control after the fact.

2 YOUR KEY DECISIONS (filled)

A

Small recurring re-purchase

STAKE

Pay within the cap, authorized categories

IMPACT

Reversible, low stake

DEADLINE

Continuous

AUTONOMY

Autonomous (within the cap)

B

Significant purchase

STAKE

Commit a high amount

IMPACT

Human confirmation required

DEADLINE

One-off

AUTONOMY

Co-decision + confirmation

C

Large or irreversible purchase

STAKE

Non-cancellable spend, outside the mandate

IMPACT

Commits the principal's money

DEADLINE

-

AUTONOMY

NON-DELEGABLE

Reading: autonomy is graded by the stake. Strong on small re-purchases within the cap, bounded on large ones, none on the irreversible without an explicit mandate. It is the mandate that makes delegation possible **in confidence**.



OBJECTIVE

Sort purchases by **criticality**, never by amount alone. It is the key tool of the case: it separates what an agent can commit on its own from what requires human confirmation.

1 THE FORMULA

$$\text{Criticality} = (\text{Impact} \times \text{Frequency} \times \text{Irreversibility}) \div 10$$

Thresholds: < 15 **low** · 15-25 **moderate** · > 25 **high**

2 THE RATED DECISIONS (filled)

Type	I	F	Ir	Criticality	Zone / autonomy
Small recurring re-purchase	2	9	1	1.8	Low · Autonomous within the cap
Subscription re-purchase	3	8	2	4.8	Low · Autonomous within the cap
Standard purchase in categories	4	7	3	8.4	Low · Supervised auto.
Purchase near the cap	6	5	5	15	Moderate · Confirmation
Significant purchase	8	4	6	19.2	Moderate · Human confirmation
Large or irreversible purchase	9	4	8	28.8	High · Non-delegable (explicit mandate)

Routing by **amount** alone would miss the nuance. Routing by **criticality** isolates the large or irreversible purchase (28.8) as non-delegable without an explicit mandate: exactly what scoped tokens reserve for the human.



OBJECTIVE Define the agent's constitution: identity, standing principles and forbidden zones. Card filled with public case data.

1 AGENT IDENTITY & MISSION

AGENT NAME

Buyer agent (scoped payment token)

ACCOUNTABLE (DDAO)

The principal user who grants the mandate

PRIMARY MISSION

Pay on the user's behalf, within the limits of a defined mandate

SCOPE

Purchases within the cap and authorized categories; excluding large purchases and the irreversible without confirmation

2 CORE PRINCIPLES

- 1 Never exceed the cap or step outside authorized categories, technically
- 2 Require human confirmation beyond a defined threshold
- 3 Tie each transaction to a bounded mandate (object, amount, duration)
- 4 Reserve the irreversible purchase for an explicit user mandate

EXAMPLE PRINCIPLES

- Respect the legal and sector framework
- Protect data and confidentiality
- Stay transparent, every decision traceable
- When in doubt, escalate rather than decide

3 FORBIDDEN ZONES (ABSOLUTE)

- 1 Commit spending beyond the mandate cap
- 2 Pay outside authorized categories or merchants
- 3 Make an irreversible purchase without human confirmation

EXAMPLE PROHIBITIONS

- X Commit spending above a threshold without validation
- X Share personal data with a third party
- X Change its own rules or scope on its own

Implicit constitution of the setup: caps, categories, duration and confirmation threshold are set **before** any action. The bound makes overrun technically impossible, not merely discouraged.



OBJECTIVE

Describe the bounds that make drift impossible beyond a threshold, like a guardrail by design. Filled with the triggers of the case.

1 TRIGGERS

Signal	Threshold	Action
Purchase amount	Beyond the mandate cap	Transaction blocked at the token level
Category or merchant	Outside the authorized list	Payment declined
Large or irreversible purchase	Beyond the confirmation threshold	Escalation, human confirmation required

2 PROCEDURE

- 1 Who decides the stop: the control at the token level, under the user's mandate
- 2 How: cap, allow-list and confirmation enforced technically
- 3 Regular testing of the bounds on borderline purchases
- 4 Log every block and every transaction (ACF-08)

Unlike Replit, the bound here is **technical**: the agent cannot exceed the cap or step outside the categories. Trust rests not on the agent's good behavior, but on the design.

Disclaimer

This study is carried out **exclusively from public information**. It constitutes **neither an audit, nor a certification, nor an official evaluation** of any actor cited. It applies the ACF framework for pedagogical and illustrative purposes. No internal, confidential, or non-public data was used. The trademarks and institutions cited belong to their respective owners.

Sources

- Visa, *Visa Intelligent Commerce* (2025 announcements)
- Mastercard, *Agent Pay* (2025 announcements)
- Google, *AP2 (Agent Payments Protocol)*; press pickups (Electronic Payments International, Digital Commerce 360, 2025)

ACF Case Study ACF-CS-019 · v1 · working document, to be reviewed before any public distribution.