Payments MAVEN™

Agentic Commerce and the Governance Gap

What happens when AI agents start moving money

Marcia Klingensmith
The Instant Payments Maven™
Smart Cities — Raleigh-Durham | July 2, 2026

Payments MAVEN™

Adoption is outpacing governance. The gap is now measurable.

77%	of federal leaders say oversight frameworks are essential, yet **fewer than 1/3** have actually implemented them	*Market Connections / ServiceNow, 2026*
20%	of federal agencies have defined policies for pre-deployment testing of agentic AI, only 8% have an incident response framework	*Market Connections / ServiceNow, 2026*
40%	of enterprise applications will include embedded AI agents by end of 2026, up from less than 5% in 2024	*Gartner, 2025*

The opportunities are real. The governance infrastructure to manage responsible execution cannot keep up.

Payments MAVEN™

Raleigh proves the pattern is already live.

Ask Raleigh and ***Daisy*** handle resident service requests today and Raleigh's first step toward a full Customer Experience Center by 2027.

Raleigh's flood early warning system already combines prediction modeling, live sensor and remote control and can remotely lower Lake Johnson before a major storm.

*The **inform → recommend → act** pattern is already running inside a regulated government environment. The governance layer underneath it has not been built yet.*

Payments MAVEN™

Every agent stack follows the same arc: inform, then recommend, then act.

Ask Raleigh & Daisy	inform → recommend → book / pay
Flood early warning	detect → recommend → dispatch / procure
Enterprise commerce agents	assess limits → apply preferences→ execute transaction

Payments MAVEN™

Agents create intent. Very little governs what happens next.

- Once an agent can recommend, a natural evolution is to expect it to act
- Acting means initiating a transaction: a payment, a deposit, a disbursement.
- Agents can now initiate payment transactions The controls and governance to support them haven’t kept pace.

Payments MAVEN™

Agents need a Real-Time Control Layer™, not just access to APIs and tools.

- Who is the user, and who authorized the agent to act on their behalf
- What is the agent allowed to do, and up to what dollar limit
- What data and rule authorized this specific action
- Is the transaction reversible, and who is liable if the agent is wrong

These are the same questions NIST's AI Risk Management Framework and the FY2026 NDAA's autonomous-systems provisions already ask.

Payments MAVEN™

“Prevention is cheaper and faster than remediation. The alternative is learning from a very expensive mistake.”

Payments MAVEN™

Engineers are solving the agent layer. Few have thought about the control layer that should reside underneath.

Major platforms are racing to define agent-to-agent and agentic commerce standards.

Engineering teams are scaling orchestration, identity, and tool access faster than settlement controls.

The unresolved question isn't “can the agent act.” It's “who's accountable when it moves money”.

Liability for agent-initiated transactions remains unsettled across every jurisdiction watching this space.

Payments MAVEN™

Once the control layer exists, the next question is what money should actually move.

Tokenized Deposits

A digital claim on a bank deposit, built for instant, programmable transfer. Stays inside the existing banking system.

Stablecoins

Broader reach, but when an agent moves money via stablecoin, explaining who authorized it, where it went, and why becomes significantly harder.

Both are programmable. The difference is where each sits on the regulated-to-flexible spectrum.

Payments MAVEN™

Each payment type fits a different risk profile.

Tokenized deposits

- Regulated B2B settlement
- Conditional, auditable disbursements
- Identity-bound, reversible transactions

Stablecoins

- Cross-border or 24/7 settlement
- Composability across agent ecosystems
- Traceability gaps create accountability exposure in auditable environments

The instrument decision should follow the use case, not the other way around.

Payments MAVEN™

“The payment rail is not the starting point. The starting point is the agentic workflow.”

Once an agent can recommend, reserve, or dispatch, the system needs a controlled money-movement layer underneath it.

Payments MAVEN™

Before any agent moves money, five questions need answers, and regulators are already asking them.

1. Who authorized the agent, and what are its limits
2. How is consent captured for each transaction
3. Is the transaction reversible, and what's the dispute process
4. Who is liable: the platform, the agent provider, or the operator
5. Can the action be tested, audited, and rapidly revalidated

Approval, oversight, testing, human involvement, auditability, revalidation — the FY2026 NDAA names these for autonomous systems. They apply just as directly to autonomous payments.

Payments MAVEN™

None of this requires waiting for the technology to mature.

It requires deciding the governance now.

- The agentic layer is being built today, in cities, in platforms, in every stack in this room
- The control layer is the piece still missing almost everywhere
- Teams that answer these questions first will deploy faster, not slower

Payments
MAVEN™

The agent decides.
The payment instrument settles.
The control layer governs.

Next step: ask whoever owns your agent stack what control layer sits beneath it.

Stay close to what's next.

Subscribe to The Instant Edge for frameworks on instant payments, AI governance, and banking modernization.

instantpaymentsmaven.substack.com | linkedin.com/in/marcekmba