BUILDING THE AGENT-READY PAYMENTS ENTERPRISE

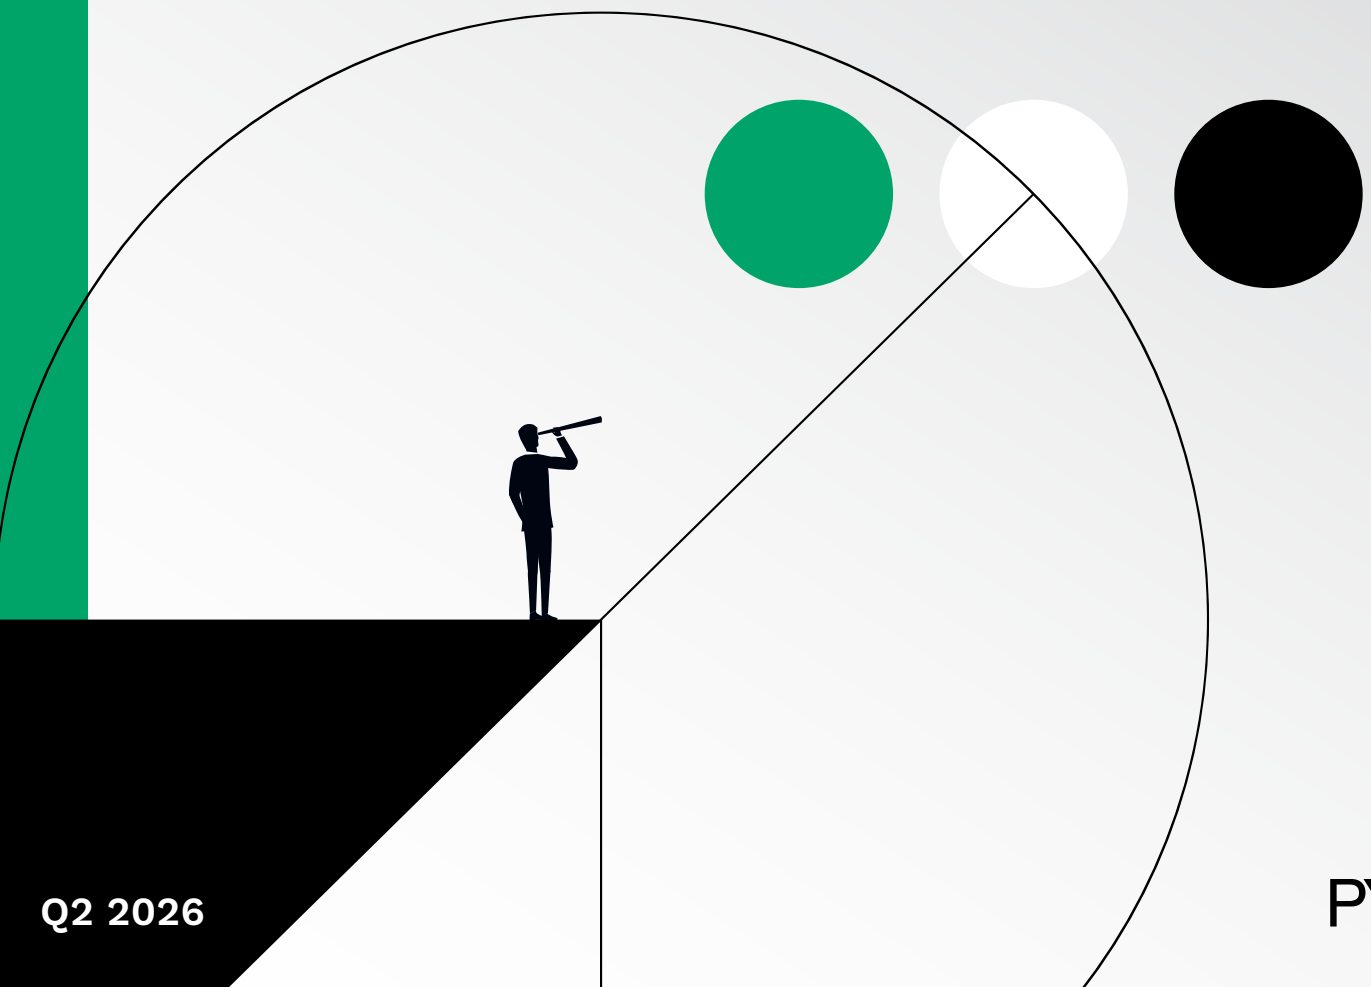

Q2 2026

PYMNTS®

TABLE OF CONTENTS

PYMNTS®

BUILDING THE AGENT-READY PAYMENTS ENTERPRISE

PYMNTS®

THE AGENTIC ENTERPRISE MEETS **THE MOVEMENT OF MONEY**

Artificial intelligence is crossing an important line inside the enterprise. It is moving from generating answers and assisting employees to executing work, coordinating systems and making decisions within defined boundaries.

That transition is already underway across payments, banking, credit, compliance, fraud, customer service and corporate finance. In this eBook, executives from Visa, FIS, Synchrony, WEX, Billtrust, i2c, Thales, Velera, Bottomline and other industry leaders describe what they are learning as AI agents move from demonstrations into real operating environments.

Their contributions reveal a broad agreement: The agentic enterprise will not be built by deploying the largest number of agents or pursuing autonomy for its own sake. It will be built by redesigning the systems, processes and decision rights around them.

That work starts with infrastructure. Agents need trusted data, connected platforms, real-time access and machine-readable rules before they can act reliably across enterprise workflows. It also requires new forms of governance. Companies must define what agents can decide, which actions require approval, when systems must escalate and who remains accountable when something goes wrong.

The essays also explore the changing role of people. As agents absorb repetitive investigation, routing, reconciliation and processing work, employees are shifting toward exception management, policy design, oversight, judgment and customer relationships. The result is less a story of wholesale replacement than one of work being redistributed between humans and machines.

Across the submissions, one principle appears repeatedly: Autonomy must be earned. Low-risk, structured tasks may be delegated first. Higher-impact decisions involving money, regulation, fraud, credit or customer trust will continue to require stronger controls and human intervention.

The agentic enterprise is therefore not a distant vision of machines running companies on their own. It is an operating model taking shape now, one workflow, permission and decision at a time.

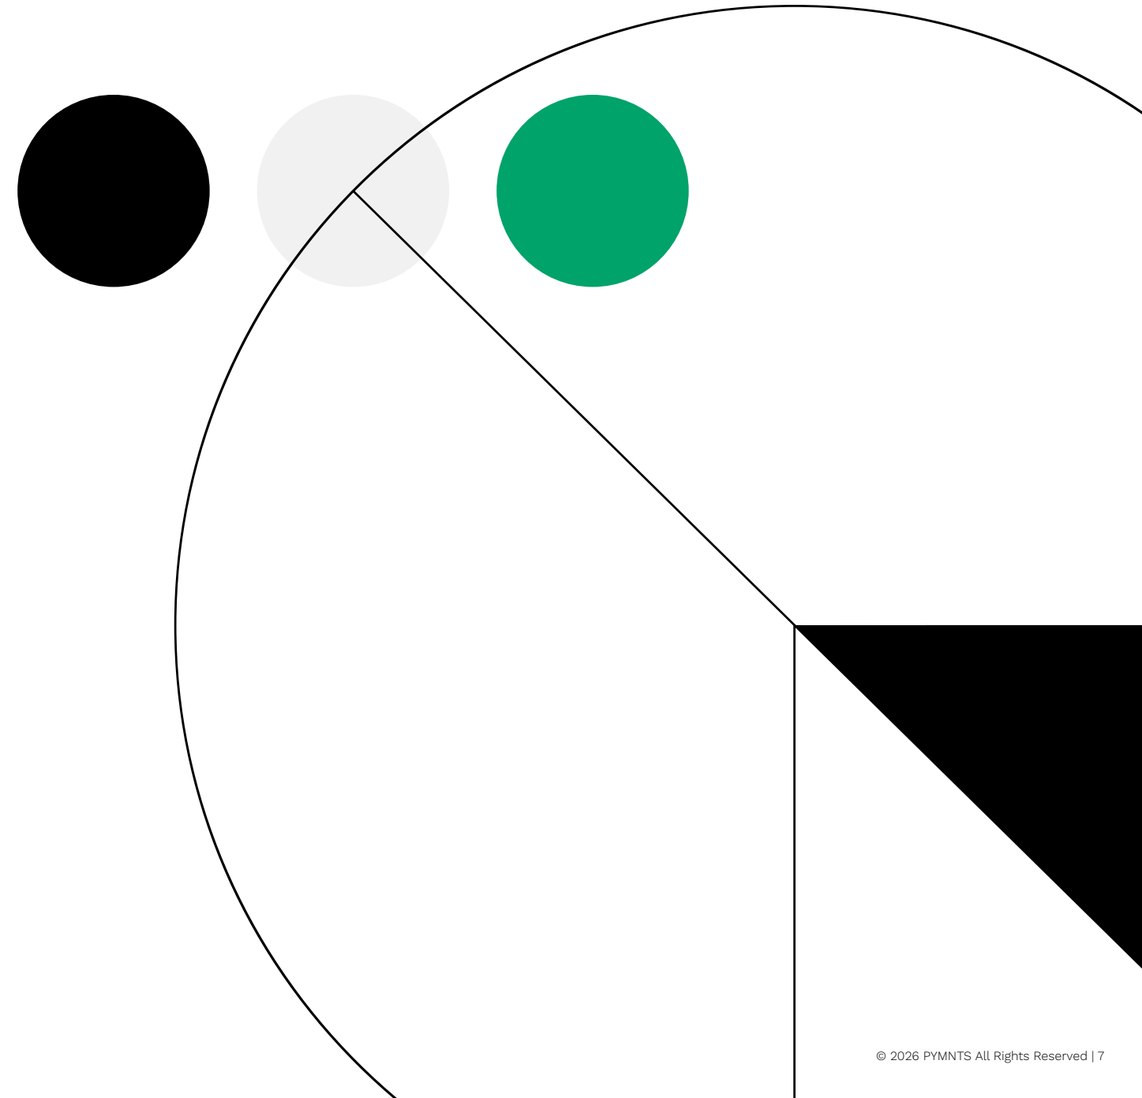

FROM COPILOTS TO COLLABORATORS: WHAT IT ACTUALLY TAKES TO BUILD AN AGENTIC ENTERPRISE

billtrust

JOHN
LANDY
Chief Technology Officer

The artificial intelligence terminology has shifted faster than most enterprise teams can adapt. Copilots became agents. Chatbots became autonomous workflows. And “agentic AI” went from a Gartner analyst term to a CFO line item in about 18 months. But the real question hasn’t changed: Are you actually getting work done differently, or just describing the same automation with better branding?

Here's what I've come to believe: The companies becoming truly agentic aren't the ones deploying the most agents. They're the ones that got serious about the infrastructure underneath them first.

THE DATA LAYER IS THE FOUNDATION

Every company has a data moat: Years of transactional history, customer behavior and operational signals. But volume alone doesn't create leverage. AI doesn't just need access to data; it needs context, metadata and domain knowledge layered on top. What's the point of terabytes of records if the agent cannot distinguish a disputed invoice from a payment preference issue?

Most agentic deployments hit this wall first. The model is capable. Tooling is available. But the underlying data is messy, siloed or ungoverned. Bad data governance means agents make flawed decisions at scale, a debt you pay later in rework, model drift and customer friction.

A meaningful step forward is the Model Context Protocol, an emerging standard that lets AI tools query live enterprise data in structured, governed ways. We've built this into our platform: A CFO can ask their AI to summarize AR risk going into quarter end and get an answer from live data in seconds, without logging into a separate system. The AI gets access to the right things, permissioned appropriately — but not everything. That distinction matters enormously as enterprises scale agent access across functions.

WHERE AGENTS ARE WORKING AND WHERE THEY'RE NOT

Real impact is happening in bounded workflows: Behavioral segmentation for collections outreach, payment policy optimization, anomaly detection in cash application. This works because the data is structured, the objective is clear and a wrong decision is recoverable. The next step is agents that don't just answer questions but take action, sending outreach, applying payments, escalating disputes, all governed by role-based permissions and a full audit trail.

The toughest lesson is that full autonomy isn't the right goal for most enterprise workflows right now. The failure mode isn't artificial intelligence that can't perform; it's AI that performs without sufficient human visibility. When an agent makes 10,000 decisions a day, a 1% error rate is a 100-problem-per-day operation. The most mature teams design for "human-on-the-loop": The agent executes, the human monitors exceptions and approves high-stakes actions. Decision rights haven't disappeared; they've been restructured.

THE ENTERPRISE FABRIC QUESTION

Becoming an agentic enterprise is ultimately an organizational decision, not a technology one. It requires agreement on what agents are authorized to do, what data they can act on, how performance is measured and how governance extends to third-party partners in your AI stack. You can't govern in the dark, and you can't govern alone.

The companies that will lead aren't the ones with the most ambitious agent roadmaps. They're the ones building the right foundation: Clean data, strong governance, meaningful human oversight and the clarity to know where autonomy earns its place.

PYMNTS®

BUILDING THE
AGENT-READY
PAYMENTS
ENTERPRISE

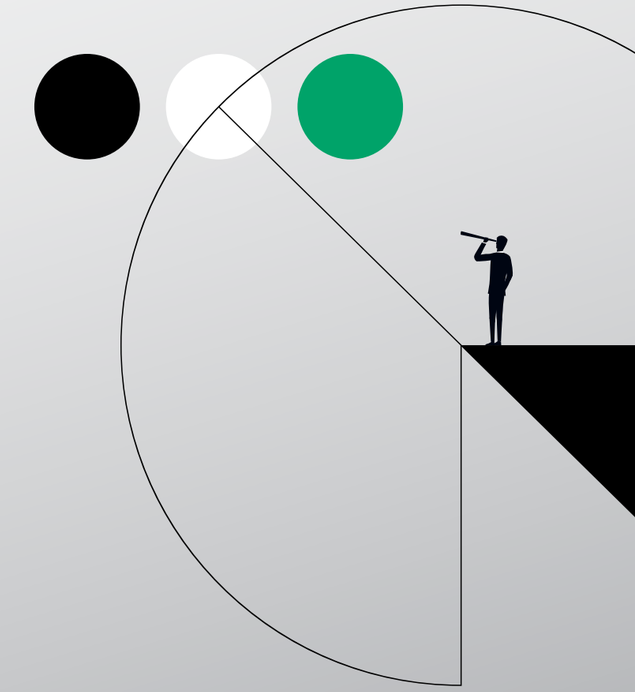

boost
payment solutions®

ILLYA
SHELL
Chief Operating Officer

AGENTIC PAYMENTS START WITH **THE RIGHT FOUNDATION**

As artificial intelligence moves from copilots to agents, the opportunity to evolve how businesses function is becoming more practical and immediate. Enterprises are starting to apply AI to real workflows where speed, accuracy and consistency matter. The bigger question is not simply what AI can do, it is what business foundation AI is being built upon.

AI agents have begun to autonomously execute defined work across systems and decision points. But in enterprises, agents do not become valuable in isolation. They become valuable when connected to trusted data, strong technology, established workflows and clear governance.

At Boost, we do not view AI as something being applied to a blank slate. We view it as an accelerator of foundational capabilities that already exist: Years of proprietary payments data, deep B2B payments expertise, established customer and partner relationships and a best-in-class technology platform built around complex enterprise payment workflows.

That foundation is what makes AI more powerful. Agentic AI grounded in proprietary data and domain knowledge can help validate information, flag exceptions and support better decisions. In the payments industry, that context is critical. Understanding how enterprise buyers, suppliers, issuers and processors interact is not something that can be recreated through generic automation alone.

Our business is already seeing this transition take shape. AI is being applied in areas such as payments processing, quality control, product design and development, client onboarding and customer service, where agents can help reduce manual effort and improve operational consistency. More importantly, agents will help us build more intelligent customer-facing solutions and experiences over time.

But agentic AI also requires discipline. Thoughtful design and planning are necessary ingredients to properly build an agentic enterprise. In payments, speed only matters if accuracy, control and trust are preserved. There is very little room for ambiguity and no room for error. That is why the future of enterprise AI cannot be unrestricted autonomy. It has to include intelligent orchestration.

In practice, AI can automate frontline processing, with another AI-enabled layer established to validate the output. Human quality control can remain in place to address issues that require judgment. The goal is not to remove people from the process, but to give them better systems, deeper insight and more technical leverage.

This is where many companies will find that autonomy is harder to scale than expected. The challenge is not always the model itself, but instead the enterprise context around the model. AI needs reliable data, mature workflows and business rules that reflect how the company operates. The agentic enterprise will be defined by companies that combine AI with proprietary assets, operational expertise and disciplined governance to generate business outcomes that could not be achieved by humans alone.

For customers, the impact should be tangible: Faster support, smoother onboarding, stronger knowledge access, more engaging solutions and payment experiences that feel less fragmented.

AI will not replace the foundation of a strong enterprise; it will amplify it. The companies best positioned for the agentic future are the ones starting with something unique to build upon.

Bottomline™

COLIN
SWAIN

Global Head of Product, Corporate Solutions

BUILDING THE AGENTIC ENTERPRISE REQUIRES MORE THAN **INTELLIGENCE, IT REQUIRES CONTROL**

Moving from manual work and processes to copilots and then to agents means a sea change in how work gets done. Now businesses can assign real responsibility to systems that can execute tasks, make decisions and operate across enterprise workflows. The promise is powerful, but the risk is real. For most companies, the challenge is no longer whether to adopt artificial intelligence (AI), but how to do so without losing control of key workflows, key decisions and creating additional risk on the whole business.

In finance, that tension is especially clear. AI has already proven its value in analysis, forecasting and reporting. But the real opportunity sits in processes like collections, payment screening and workflow, cash application and risk management. These are the workflows that directly impact cash flow, working capital and managing financial risk. They are also the areas where autonomy is hardest to scale, because the cost of errors is so high.

What separates early experimentation from an agentic enterprise is governance. Assigning work to agents and agentic workflows requires clear rules, defined decision rights and full visibility into how and why actions are taken. Without that, AI becomes another layer of risk rather than a source of value. This is where many organizations are learning that autonomy and agentic AI is harder to implement than expected. This is less because the models are insufficient — though they are improving all the time — but because enterprise systems, data, and controls were not built to support the shift to AI.

The path forward, then, is controlled autonomy rather than full autonomy.

Leading organizations are embedding AI directly into workflows, rather than layering the technology on top of core processes. In this model, agents do not operate independently, but within defined processes that are guided by policies, thresholds and approvals set by the business. Every action is traceable, auditable and aligned to financial controls. This shifts AI from an insight-delivering tool to something that can effectively execute the work that must be done. In this scenario, AI can impact the metrics that matter.

We see this progression clearly through the CFO AI Maturity Model. Most organizations start with assistive AI, where technology helps individuals complete tasks more efficiently. From there, workflows become more automated but still require human intervention at key points. The next phase introduces agentic execution, where systems can resolve exceptions, take actions and operate within defined guardrails. Ultimately, the goal is outcome-driven finance, where AI continuously optimizes key metrics like cash flow, risk exposure and operational efficiency.

At each stage, the role of the human employee changes. The focus shifts from doing the day-to-day work to defining the rules, monitoring outcomes and controlling performance. Finance functions should think of this as elevating the roles employees play, rather than removing people from the process.

The companies making progress are not chasing autonomy for its own sake, to check a box or to send out a splashy press release. These businesses are redesigning workflows starting with high-value use cases and scaling incrementally. They are connecting data across systems, embedding AI where decisions are made, and ensuring that governance is built into workflows from the start.

The agentic enterprise will not be defined by how much work AI can get done. Instead, it will be defined by how much confidence organizations can place in AI to achieve desired outcomes without increasing risk.

MLADEN
VLADIC

Head of Product Management, Payment Networks

BUILDING THE AGENTIC ENTERPRISE: LESSONS FROM THE PAYMENT MOMENT

As artificial intelligence evolves from copilots to agents that execute work across enterprise systems, what does it actually take to make autonomy real? In payments and commerce, we're already seeing the answer take shape at the point where value is created: The purchase.

WHERE AGENTS ARE RESHAPING WORKFLOWS AND DECISION RIGHTS

Agentic commerce has arrived — it's here already. A new generation of AI agents are emerging that now search, source, negotiate and complete purchases on our behalf (consumers), transacting at superhuman speed within rules their owners define. McKinsey projects these agents could help generate up to $1 trillion in orchestrated retail revenue in the U.S. alone by 2030, with global projections reaching as high as $5 trillion.

This shift redistributes decision rights. When an agent chooses where to shop and how to pay, loyalty, offers and benefits become part of its decision logic, not a post-purchase afterthought. The rise of item-level insights has been a game changer. This shift from transactional to relational payments was a trend we saw coming at FIS, and it has been great to see our early investments in advanced tech pay off.

That distinction between transaction-level and item-level intelligence matters. Transaction signals tell you where someone shopped and how much they spent. Item-level intelligence gives you a digital-receipt view, line by line, so value can be applied with item-level precision.

TRANSFORMING THE CUSTOMER EXPERIENCE

The real transformation is timing. When value shows up at the moment of purchase, it feels less like points accounting and more like recognition. FIS Smart Basket, a modular, funder-agnostic network, brings offers, payments and loyalty into a single infrastructure for banks, credit unions, issuers, merchants and consumer brands. With an advanced scoring engine and real-time item-level intelligence, customers earn and enjoy rewards instantly, eliminating the traditional lag between the two.

RISK CONTROLS AND GOVERNANCE GUARDRAILS

Autonomy without guardrails is a liability. The winners will be the ecosystems that securely embed intelligence into payments so agents can act with guardrails, relevance and trust. Three principles guide our approach:

- Agents use only approved payment methods, so humans stay in control.
- Transactions execute within existing authorization, authentication and dispute networks, keeping financial institutions integral to security and oversight.
- Spending limits, permissions and rules of engagement stay with the account holder and the institution, who can adjust or revoke an agent's authority at any time.

WHERE AUTONOMY IS HARDER TO SCALE THAN EXPECTED

Here's the honest part. Scaling agentic systems is tougher than the headlines suggest. The most scalable ecosystems start with what's universally available, transaction-level signals across the merchant landscape, then deepen into item-level precision where item visibility exists.

Two barriers slow adoption. First, fragmented infrastructure. Card issuers, sellers and brands must collaborate within a unified, single infrastructure, something most enterprises haven't built. Second, first impressions: If a customer's first interaction with this experience isn't ideal, they won't come back to it quickly.

The agentic enterprise won't be built through futuristic autonomy. It will be built through deliberate decisions about where to trust agents, where humans stay in control and which workflows to redesign first. The enterprises that win won't be the ones running the most agents — they'll be the ones that own the connective layer agents act through, with the institution still setting the rules. In payments, that work is already underway.

PYMNTS®

BUILDING THE
AGENT-READY PAYMENTS ENTERPRISE

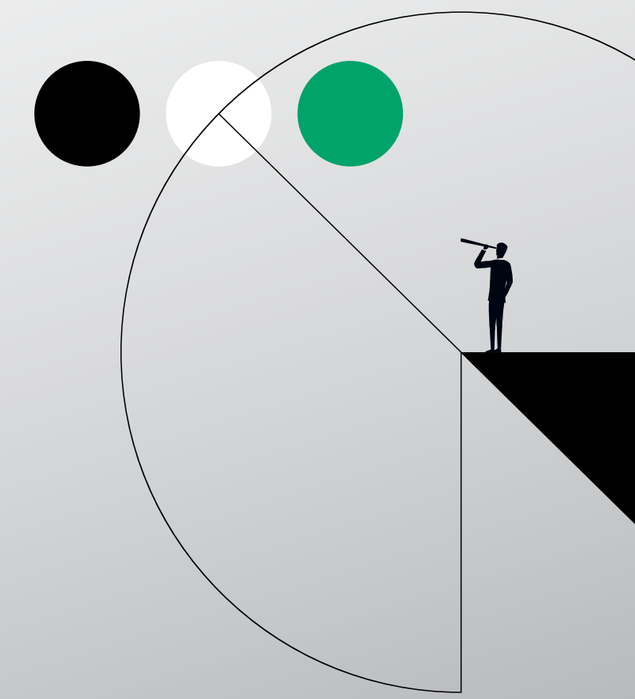

Flagright

MADHU G.
NADIG

Co-Founder and Chief Technology Officer

THE AGENTIC ENTERPRISE WILL BE BUILT ON DECISION RIGHTS, **NOT MODEL INTELLIGENCE**

The agentic enterprise will not be defined by how many tasks AI can perform. It will be defined by how safely an organization can delegate authority.

That distinction matters. Copilots advise; agents act. Once an artificial intelligence system can retrieve data, update records, resolve cases, communicate with customers or trigger down-stream workflows, the central question is no longer model intelligence. It is decision design: What the agent may decide, what evidence it must collect, when it must escalate and who remains accountable.

The most scalable model will be bounded autonomy. Enterprises should give agents narrow mandates inside explicit operating rules, not open-ended permission to "handle" a function. A well-designed agent should execute an approved playbook, document each step, expose the evidence behind its conclusion and stop when confidence, policy or risk thresholds are breached.

Financial crime compliance shows where this is already practical. Agents can investigate routine monitoring and screening alerts, gather information across systems, compare activity against institutional policy, generate narratives and resolve clear false positives. Human investigators can then focus on complex networks, ambiguous behavior, regulatory judgment and customer decisions. The value is not simply faster analysis. It is a redesigned operating model in which repetitive evidence collection is automated while consequential judgment remains deliberate.

This changes staffing, but not through indiscriminate headcount removal. Entry-level work will shift away from manual triage. Teams will need more people who can translate policy into executable workflows, test agent behavior, monitor exceptions and improve controls. Domain expertise becomes more valuable because agents scale the quality and the weakness of the instructions they receive.

This also changes economics. Functions that once scaled by adding reviewers can absorb higher volumes without matching increases in capacity. But the resulting savings should fund stronger investigation, control design and exception handling, not erase institutional accountability.

Customer experience improves for the same reason. When low-risk cases are resolved quickly and consistently, legitimate customers face fewer delays. But an agent should not optimize convenience at the expense of defensibility. In regulated or high-impact workflows, speed without traceability creates a larger liability, not a better experience.

Most autonomy programs will stall because enterprises automate tasks before fixing the systems around them. Fragmented data, contradictory procedures, unclear decision ownership and weak audit trails cannot be solved by a more capable model. Agents amplify operational design. If the underlying process is incoherent, autonomy makes that incoherence faster.

Building an agentic enterprise therefore requires four foundations: Machine-readable policies, governed access to trusted data, pre-deployment testing against historical scenarios, and continuous observability after deployment. Every action should be attributable to a policy version, model version, data source and defined authority level. Human override must be operational, not ceremonial.

My prediction is that leading enterprises will stop measuring AI progress by the number of agents deployed. They will measure the percentage of decisions that can be delegated with explicit controls, reliable evidence and accountable escalation. The winners will not be the companies that pursue maximum autonomy. They will be the ones that make autonomy governable.

AMIR
WAIN
CEO

YOUR AGENTIC AI STRATEGY IS ONLY AS STRONG AS THE **FOUNDATION BENEATH IT**

I think most companies building agentic AI are answering the wrong question. They ask which model to deploy. The harder question is whether the infrastructure underneath was ever designed to let something act on its own.

I've watched AI move through three eras: Rule-based systems in the 1980s, machine learning that detected patterns starting around 2010, and now agentic AI — systems that reason, act autonomously, orchestrate across systems, and learn from every outcome. Generative AI gave us reasoning. Agentic AI adds perception, action and memory.

The companies winning in this era aren't necessarily the ones with the most sophisticated artificial intelligence strategy. They're the ones whose architecture was built to support it. An agent can only orchestrate across systems it can see and only act in real time if the underlying platform responds to events as they happen. When AI is layered onto fragmented systems and disconnected data, what looks like an AI initiative quickly becomes a data integration project.

We learned this through experience. Decisions we made years ago to build a unified data model across our platform are paying off in ways I didn't fully anticipate. Our fraud systems can read signals across a customer's entire relationship, not just a single transaction type, and learn continuously from outcomes. That foundation allows us to move decisions closer to the transaction itself.

The clearest example is fraud management. What once required an analyst review and a call-center process can increasingly be handled within the transaction flow. An anomaly is detected, action is initiated and the customer experience becomes a quick confirmation rather than an embarrassing decline or a lengthy dispute process. That's not automating a step. It's removing the step entirely.

As autonomy increases, decision rights shift. Judgments that once belonged to people increasingly belong to systems, with humans positioned to intervene rather than approve every action. Getting that balance right matters. Too much oversight limits the value of autonomy. Too little creates unnecessary risk.

That's why governance becomes critical. Permissioning, audit trails and human oversight cannot be afterthoughts. The organizations scaling agentic AI successfully are building those guardrails before they need them, not after an incident exposes the gaps.

The build-versus-buy conversation is changing, too. Unless an organization operates at the scale of the largest players in its industry, building proprietary AI infrastructure from scratch is rarely the best use of capital or talent. The better question is which capabilities truly differentiate your business and which have already been solved well by trusted partners. Build where it creates competitive advantage. Partner where it doesn't.

What hasn't changed is the discipline required to evaluate any artificial intelligence use case: Economic benefit, structured data, real scale and auditability. Agentic AI doesn't lower that bar. It raises the stakes, because these systems aren't just advising anymore. They're acting.

maverick

BEN
GRIEFER

President and Chief Operating Officer

WHY AGENTIC AI DEMANDS **STRONGER RISK CONTROLS**

AI is entering a new phase. Copilots that assist employees are giving way to agents that can plan, decide and act across enterprise systems. For payments companies, that shift is no longer theoretical. It is already reshaping how customers experience service, how risk gets priced and how quickly a business can grow without growing its exposure alongside it.

Consider fraud scoring. An agent that flags a suspicious transaction and routes it for review adds speed without adding risk. An agent that autonomously declines a legitimate customer, or approves a fraudulent one, does something different: It makes a business decision on the company's behalf, in real time, at the exact moment a customer relationship is won or lost. That is the boundary every payments executive needs to understand clearly. Not every agentic capability carries the same consequence, and treating them as interchangeable is where governance breaks down.

The same pattern shows up in chargeback resolution and partner reconciliation. Agents can compile evidence, match transaction records and identify discrepancies far faster than manual review. What they should not do, in our view, is issue the final ruling on a disputed high-value chargeback or resolve a reconciliation gap with a partner bank without a person signing off. The cost of being wrong at scale, across thousands of transactions, is simply too high to fully delegate.

This is also a business model question, not just an operational one. Agentic capability changes what a company can credibly promise customers and partners: Faster onboarding, faster dispute resolution and faster underwriting decisions. At Maverick, we see this value as AI helps accelerate approvals, surface the right processor and portfolio data to our teams and improve decision-making. As a full-service payments provider, we use these capabilities to help us connect merchants with the best-fit solution more quickly while maintaining the human oversight and accountability that complex payments decisions require. But promises made possible by AI still have to be kept accountable to a person when something goes wrong. Customers do not care whether an agent or an employee made the call. They care that someone owns the outcome. That expectation does not change just because the technology got faster.

The organizations getting this right are rebuilding decision rights around risk, not around convenience. They are asking which workflows can tolerate full automation, which require a human checkpoint and which should never be fully autonomous regardless of how sophisticated the underlying model becomes. They are also being transparent internally and externally about where third-party models and partner platforms sit inside those workflows, because accountability does not transfer just because the technology was licensed rather than built.

None of this slows a company down if governance is designed in from the start rather than bolted on afterward. The companies that will lead in agentic payments are not the ones deploying the most agents. They are the ones that know precisely which decisions their agents are allowed to make, which ones still require a person, and why.

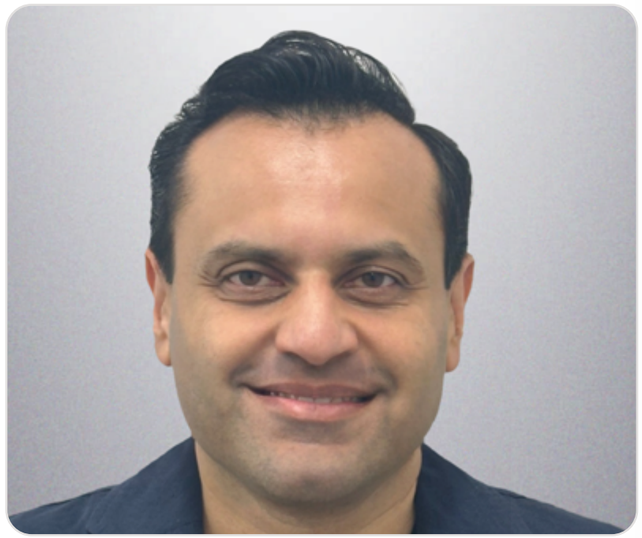

splitit

NANDAN
SHETH
CEO

THE PAYMENT MODEL BUILT FOR **AI AGENTS**

Building an agentic enterprise means rethinking how you do business. As artificial intelligence (AI) agents start making decisions, enterprises will have to determine not just what agents can do, but also which processes must change to allow agents to function properly.

In payments, that means companies must rethink which parts of their payment model breaks down because human intervention is required.

Traditional buy now, pay later (BNPL) is one of those areas. Over the past 10 years, BNPL has helped retailers solve the affordability problem for countless customers. And by most measures it raises conversion and average order value.

Nevertheless, BNPL was built around a moment that stops AI agents cold, and the companies still waiting on that moment are going to find themselves waiting alone.

In a BNPL transaction, a consumer applies for short-term credit at the point of purchase, receives a real-time approval decision and completes the transaction through a provider like Affirm. This model was created for a human shopper waiting in front of a screen. If there's a problem, he or she can simply fill out a form, accept the terms and be redirected to a third-party flow.

But now, in the emerging world of agentic commerce, AI agents act on a shopper's behalf to find, evaluate and complete a purchase. There is no human in the loop to handle a new credit application. The transaction either completes or it doesn't. With new-credit BNPL approval rates running between 35% and 40%, 6 in 10 BNPL transactions handled by AI could fail at the payment step — and the worst part is, the retailer won't even know. The agent will just move on to another merchant.

Businesses deploying agents might only discover this limitation in production when the cost of rearchitecting a payment integration is significantly higher than the cost of modeling it correctly from the start.

One emerging solution — installment payments linked to a consumer's existing credit card — provides the certainty that BNPL lacks. Here's how it works: The merchant authorizes the full purchase amount against the shopper's existing credit card at checkout, then splits that authorized amount into monthly payments. That means no new credit application or third-party loan, just installments drawn against credit the shopper already has.

Card-linked installments help prevent payment failures because the credit decision is made when the card is issued. The agent is drawing on available credit, not triggering a real-time lending decision.

Just as important, card-linked models select consumers with established credit and available capacity. These consumers typically make higher-AOV purchases and are repeat buyers. Card-linked BNPL is a high-quality, low-risk growth engine.

In the world of agentic commerce, checkout is where the hard work begins. Retailers who get it right will be the ones agents keep coming back to.

synchrony

MIKE
STORIALE

SVP, AI Technology and Transformation

AN AGENTIC-READY ENTERPRISE **STARTS WITH TRUST**

Building an agentic-ready enterprise starts with trust between a company and its employees. As our CEO Brian Doubles has said: “Our culture is built on trust — giving our people the freedom to innovate, experiment and apply technologies like AI in ways that make a real difference for our customers and partners.”

At Synchrony, years of consistent listening, collaboration and follow-through have ensured that employees trust that we’re using artificial intelligence to make their work better.

When they told us they wanted more access, training and use cases of how to use AI tools daily, we rolled out SynchronyGPT across the company, now paired with an AI Field Guide covering best practices and prompt training so employees could immediately apply it to their everyday work.

Adoption is strong, with nearly 100% of our professional workforce using AI tools and 80% of employees believing AI will positively impact their careers. Importantly, employee trust is high — 90% of employees trust Synchrony to use AI fairly, ethically and responsibly.

We see AI as additive, where people remain at the core as the creative idea leads, critical thinkers, and decisionmakers. Artificial intelligence frees our best talent from tactical work and redeploys them to more meaningful, innovative work.

Agentic AI also requires transforming the relationship between companies and customers. As purchasing authority shifts to AI acting under permission, hard questions must be answered. Who is authorized? What actions can an agent take? How is consent captured and enforced? Who is accountable for disputes, fraud or errors? Companies that answer those questions clearly and perform consistently will be the safest choice for consumers to delegate to and the easiest choice for agents to recommend.

Trust in this context becomes much more than a brand message. It is measurable performance by AI agents including predictable approvals, low fraud loss, fast and fair dispute resolution, transparent terms, strong identity protection, actual performance in customer service and clear boundaries on agent behavior. Agents will steer spend toward institutions that deliver reliable outcomes because reliability reduces risk for the consumer and friction for the merchant. When software becomes the shopper, trust becomes the last durable advantage.

How do you create AI agents that meet these high standards? First, prioritize collaboration. By aligning with frameworks like Google's Agent Payments Protocol (AP2), Universal Commerce Protocol (UCP) and Model Context Protocol (MCP), financial services companies can help shape standards for how AI agents securely find products, place orders and make payments.

Financing products must also be redesigned for a world where buyers include agents, not just consumers. Rewards, fees and financing terms need structured, machine-readable formats; fine print that diverges from what an agent ingests will send spending elsewhere. Fraud and authorization models must recognize agent patterns, and humans must remain accountable for outcomes, especially in financial services, where explainability isn't optional.

Finally, issuers should recognize that retailer agents can be natural partners. Both parties prioritize high approval rates and fast decisions. Piloting agent-aware offers with merchants positions issuers as the default choice for machine-to-machine commerce.

THALES

AMÉLIE
TOURNANT FOUREL

Head of AI-Driven Operations Platform

FROM AI ASSISTANTS TO OPERATIONAL AGENTS: WHERE ISSUERS CAN ACT NOW

As artificial intelligence shifts from copilots to agents, the focus is moving from assistance to execution. The real question is no longer how AI can support employees, but where systems can be trusted to take on real work.

Progress will not come from adding intelligence on top of existing systems. It will come from selecting the right workflows and building the operational conditions that allow AI to perform at scale. This also requires defining clear guardrails to ensure trust and control.

So where should banks start?

For issuing banks under pressure to modernize operations, improve responsiveness and reduce manual overhead, the most immediate opportunity is operational. Card issuance processes, such as card request and lifecycle management, sit at the core of customer servicing. These processes are high-volume, structured and fundamental to customer experience and brand perception. Yet they remain fragmented across systems and still rely on manual handoffs.

Take card replacement as an example. A single request can involve physical and virtual card issuance, address verification, PIN management, token updates and interactions with external partners such as personalization bureaus. Today, these steps are often disconnected, slowing execution and increasing complexity. Orchestrated end to end, they can enable faster resolution and smoother customer experience.

Another key characteristic of card issuance operations is that they rely on established data-driven workflows. This makes them strong candidates for agent-driven execution, provided banks can coordinate actions across systems, from customer request through to fulfillment. The barrier, however, is not the AI but the environment in which it operates.

Many issuers still rely on architectures not designed for real-time coordination: Batch processing, limited interoperability and rigid workflows. As a result, they struggle to trigger actions dynamically or connect execution across systems. This is why many current “AI-powered” capabilities remain limited, layering orchestration on top of legacy systems rather than actually transforming execution.

Building an agentic enterprise requires a different foundation. Before autonomy comes real-time visibility and control: Understanding what is happening across data and operations and enabling systems to act on it consistently.

For banks, this leads to a clear priority: Transform operations into connected, API-driven processes. APIs enable real-time interaction between systems, make data accessible and allow actions to be executed consistently across environments. They turn static workflows into dynamic execution layers where agents can operate effectively.

Equally important is how control is designed. Execution must follow clear rules, escalation paths and full traceability. Routine actions can be handled by systems, while humans retain authority over higher-risk situations. The objective is not to remove oversight, but to make it more targeted and effective. For sustainable trust and security, all these guardrails need to be in place from the start.

The implication is clear: Success with agentic AI will depend less on model sophistication and more on how well banks redesign their operational foundations, not only within their own systems, but across their supplier ecosystem, ensuring seamless connectivity, real-time data flows and well-governed decision-making.

The path forward is not about adding intelligence on top. It is about building operations that intelligence can run.

thredd

JIM
MCCARTHY
CEO

AI AGENTS CAN DECIDE. SOMEONE STILL HAS TO DECIDE **WHAT IS ALLOWED**

The interesting question about the agentic enterprise is not whether agents can act. They clearly can. It's whether enterprises can scale that action safely — and the evidence so far is that autonomy is harder to scale than the demos suggest. The bottleneck is not intelligence. It's permission, control and proof.

Businesses are putting policies and governance in place to safely scale generative AI across their enterprise. At Thredd, green shoots of innovation are sprouting daily, leveraging agentic capabilities for fraud, credit, sales, billing automation and client servicing. But we see the challenge, and equally the opportunity, for agentic experiences more acutely on the payment side, where an agent recommending a purchase, not all that interesting, becomes useful when it's trusted to initiate and complete a payment. At that moment the question stops being "can the agent pay?" and becomes "what is this agent permitted to do, on whose authority, under what limits, and how is that permission proven, monitored and revoked in the milliseconds an authorization takes?"

That's a payments problem before it's an AI problem, and payments has a head start most people underestimate. At every technology inflection point during my career — eCommerce, mobile commerce, the launch of Apple Pay, cryptocurrency — someone predicted the end of cards, and yet each time the trust infrastructure proved far harder to replicate than the rails were to reinvent. Tokenization, scheme rules, dispute and chargeback rights, issuers that underwrite risk and bind consumers to credentials — these aren't legacy baggage. They're the exact controls agentic commerce needs, and they already exist.

But raw materials aren't readiness, and this is where the conversation gets too casual. A network token already carries merchant and category restrictions and can be revoked providing a real head start, not an answer on its own. The new work is binding a verifiable mandate to that credential which is essentially a cryptographic proof of what the consumer authorized their agent to do, and then surfacing, at authorization, a signal that an agent rather than a human initiated the transaction, so the issuer can decide in real time rather than discover it after the fact. Visa and Mastercard are standardizing this through Intelligent Commerce and Agent Pay respectively; we monitor and, in some cases, build those frameworks while staying deliberately neutral on the merchant-side protocols, so our clients, issuers and program managers, are insulated from bets that haven't been resolved.

It's also why a connection standard, on its own, doesn't solve this. Protocols like MCP matter because they standardize how agents plug into systems. They don't answer who authorized this agent, what the consumer actually intended, whether the behavior is anomalous or who is liable when it's disputed. Those are issuer-layer questions, and they don't get easier by stacking another protocol on top of them.

Liability is the one the industry hasn't fully answered, and the one I'd tell any enterprise to watch. Fraud detection is being retuned for a new class of automated behavior. The old question was, "is this the genuine customer?" The emerging one is, "is this still what the customer authorized the agent to do?" That's necessary but not sufficient. The harder case is the agent that behaves

exactly as mandated and the cardholder disputes anyway. Today's chargeback rules were never written for a third party acting on standing instructions. My view is that verifiable intent, proven at authorization, must become the basis for reallocating that liability — and the issuers who can prove intent will be the ones who can actually underwrite agentic commerce, not merely permit it.

For a card issuer, the point of sale is the tip of the agentic iceberg. Strip away the consumer experience and the work of card issuing is repetitive, rules-based, operational work. Things like reconciling settlement files, onboarding and screening customers, assembling dispute evidence, servicing cardholders and producing regulatory reports represent high-headcount work that has scaled almost linearly with volume. That's exactly the shape of tasks agents are good at, and where costs can actually be stripped out of businesses: An agent that reconciles a file and escalates only the genuine inconsistencies, or triages an onboarding case, breaks the link between growth and headcount that has constrained this industry for years. Programmable payments extend the same logic to money movement itself where value that moves on conditions and rules set in line with the transaction rather than preset instructions configured at the program level, with an agent orchestrating inside guardrails set by the customer.

An agent reconciling a ledger, or releasing a payment on a rule, needs exactly what an agent at checkout needs: Defined permission, proof of what it was authorized to do, an audit trail and the ability to revoke it the moment something looks wrong. Solve that governance problem once and you apply it across the operation; that's the compounding advantage, and why the trust layer is worth owning. It's also where discipline matters most, because the back office is where autonomy is easiest to justify and hardest to trust. A silent reconciliation error doesn't announce itself, it compounds. Those who capture these savings will be the ones who put escalation thresholds, manual review thresholds and provable authority in place first — taking cost out rather than trading labor cost for risk.

The agentic enterprise won't be defined by how much autonomy it hands to AI, but by how well it governs that autonomy where liability shifts and value is created. The winners won't just have smarter agents. They'll have agent-ready infrastructure beneath them, verifiable intent underwriting it, and partners willing to absorb that complexity so they can move quickly.

velera

DENISE
STEVENS

EVP, Chief Product and Technology Officer

THE PEOPLE-FIRST PATH TO **AGENTIC ENTERPRISE**

Agentic artificial intelligence is beginning to reshape enterprise workflows across financial services, moving from tools that assist work to systems that can execute it. For credit unions, that shift creates an inherent tension: Autonomy can make operations faster and more consistent, but it must not weaken the human judgment and member focus at the heart of the cooperative model. The challenge is introducing agentic AI in ways that reinforce — not replace — the credit union philosophy of “people helping people.”

Building a mission-aligned agentic enterprise starts with three fundamentals:

1. IMPLEMENT A PEOPLE-FIRST OPERATING MODEL

Agentic AI cannot replace human empathy and discernment, but it can take on routine tasks such as gathering data, drafting responses and recommending next steps. This frees employees to focus on relationship management and complex problem-solving.

These human-AI partnerships require practical training and clearly defined decision rights. Employees must know when to trust, challenge or override AI actions, ensuring AI strengthens workforce effectiveness rather than replacing human judgment.

2. BUILD DECISION-READY DATA AND CONNECTED SYSTEMS

Agentic AI should not be measured by how many tasks it can complete, but by how effectively it enables actions, decisions and escalation across systems. That requires reliable, connected data.

When data is fragmented, employees must rebuild member context manually — and AI outputs become less reliable. When systems are integrated and data is structured, AI can surface timely insights, support faster responses and improve the consistency of decisions across the organization.

3. CREATE GOVERNANCE FRAMEWORKS THAT PROTECT TRUST AT SCALE

The more responsibility AI assumes, the more important governance becomes. AI agents can make mistakes, but organizations are also learning that autonomy becomes harder to scale when workflows involve exceptions, regulatory requirements or complex member circumstances. This means credit unions need clear guardrails around where AI can act independently and how outcomes are monitored over time.

Within this framework, responsible governance spans data quality, security, transparency, model oversight and clearly defined review points. Those safeguards allow credit unions to adopt AI confidently while maintaining member trust.

WHERE CREDIT UNIONS CAN START TODAY

As organizations evaluate agentic AI investments, the question is not simply what can be automated, but which applications deliver meaningful business value. That value may come through greater efficiency, faster service, more consistent decisions, stronger risk management or an improved member experience. The strongest use cases are those where the benefits clearly outweigh the cost and complexity of implementation.

Many credit unions are not yet operating fully agentic systems, but they are beginning to redesign workflows around shared decision-making between agents and employees. The strongest starting points are workflows where risk can be managed and human oversight remains central, including fraud alert triage, dispute intake, member service routing and lending preparation.

In each case, AI gathers context, structures information, recommends next actions and escalates exceptions — while humans retain final authority. These workflows help credit unions build confidence in agentic capabilities without compromising trust.

As agentic AI becomes more mainstream, credit unions don't need to be the first to implement every new technology. The strongest agentic enterprises will be those that intentionally distribute work between agents and people — scaling autonomy where it adds value, protecting decision rights, strengthening employee confidence and preserving the trust that defines the member relationship.

PYMNTS®

BUILDING THE AGENT-READY PAYMENTS ENTERPRISE

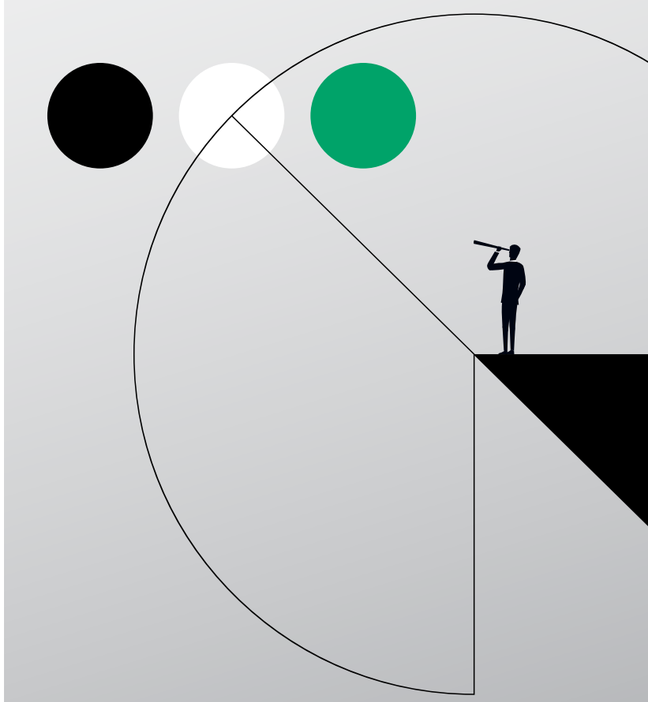

VISA

RUBAIL
BIRWADKER

SVP, Global Head of Growth and Partnerships

AI AGENTS CHANGE THE WORK, BUT **HUMANS STILL SET THE RULES**

Every major technology wave changes how we work. From the printing press to the loom, from the typewriter to the web, new technology requires new ways of working.

That's why I see building an agentic enterprise as just as much of an organizational challenge as a technology one. It's a mind-set shift that requires us to organize our work in a new way. For instance, as artificial intelligence agents take on more operational responsibilities, employees are increasingly focused on organizing how agents work and guiding decisions based on their specific domain expertise.

We're also seeing employees use AI coding assistants like Claude Code or OpenAI's Codex to build new workflows to coordinate actions across agents and create fit-for-purpose solutions that solve challenges in real time.

Many of these new tools activate existing data, unlocking fresh insights and business value. The value often comes not from a single breakthrough capability, but from lifting layers of complexity in ways that reduce manual coordination and help us move faster. The best implementations reveal new opportunities that were not visible before, tapping agents to gather information from multiple systems and connect the dots.

When it comes to true autonomy, most organizations have a ways to go. And that's a good thing, as there's quite a bit that goes into true agentic autonomy. Success requires embedding trust, transparency and accountability into AI strategies from the outset. Agentic systems need clear guardrails that limit access to an organization's data in a way that protects privacy while still ensuring that teams get the most value from their artificial intelligence tools.

Enterprises must also thoughtfully balance agentic autonomy with human oversight. It's not a binary choice, as different workflows require different levels of human involvement. Many tasks will continue to require human judgment, particularly when decisions involve financial risk, regulatory obligations, customer trust or broader business strategy.

We also must keep the end goal always in mind: Tapping both human and machine capabilities in ways that create smarter decisions, faster execution and stronger business outcomes. That's the true value of the agentic enterprise: The best of both worlds.

PYMNTS®

BUILDING THE AGENT-READY PAYMENTS ENTERPRISE

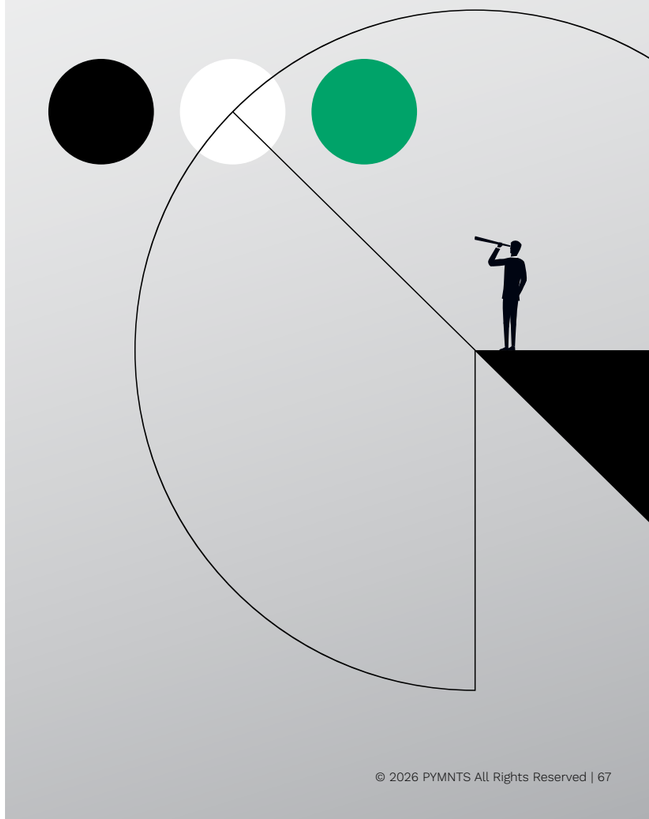

Volanté

DEEPAK
GUPTA

Chief Product, Engineering and Delivery Officer

THE AGENTIC ENTERPRISE WILL BE BUILT ON GOVERNED CONFIDENCE, **NOT BLIND AUTONOMY**

The inflection in agentic AI won't be defined by what agents can do, since that capability is arriving regardless. It will be defined by which banks can prove what their agents did. The next phase of agentic AI in payments will be shaped by a practical question: "How much work can agents take on without weakening control?"

That question guided our development of Vol360i, Volante's agentic AI solution for payments. Banks are ready for artificial intelligence to do more than summarize information or suggest next steps, but they are not asking for unchecked automation or another AI layer that creates more work for operations, risk or compliance. They need agents that can work within the payment flow, act when evidence is strong, escalate when risk is higher and show exactly how each recommendation was made.

For financial institutions, there is little to no room for trial and error. A payment cannot be fixed casually after the fact. It touches liquidity and the customer's expectation that money will move as intended. The economics now point the same way. As inference costs fall, running an agent gets cheap, but in payments, the dominant cost is the time spent on investigation and clawback, and the liquidity hit. When intelligence becomes inexpensive, the binding constraint shifts from compute to control, which makes a confidence-based operating model a rational choice.

AGENTS NEED TO LIVE IN THE WORKFLOW

A payment workflow is not the kind of process where banks can automate first and clean up the consequences later. That's why it is essential to design true value-adding solutions around agents who act where payment decisions are made.

Vol360i is built around four capabilities:

1. Sense Agents monitor for risks before they affect performance.

2. Predict Agents determine the best outcome for each payment, supporting faster, more cost-efficient decisions.

3. Prevent Agents identify potential failures before they occur, helping reduce customer-impacting errors.

4. Repair Agents resolve payment issues in real time, so operators spend less time on repetitive exception handling.

Together, these capabilities move banks away from static rules and manual queues. Further, they prepare banks for what comes next: A growing share of payments initiated by agents acting on behalf of humans. The financial institution that can govern agent-initiated payments — verifying intent, scoring confidence and escalating anomalies — will strengthen their competitive advantage as the counterparty becomes non-human.

GOVERNANCE IS NON-NEGOTIABLE

The AI race is also a governance test. Speed only matters if banks can explain and control the decisions that agents make.

Governance must shape how agents are designed, tested and monitored from the start. Banks need to understand why an agent made a recommendation and whether a human accepted, changed or rejected it. Each decision should leave a clear record, including the data behind the recommendation and the confidence level assigned to it.

That record only means something on infrastructure the bank governs — deployed in-tenant, with data kept resident and no payment information passing through shared inference. Governed confidence requires governed compute.

For DORA-regulated institutions, private AI deployment is what turns the audit trail from a promise into a control.

That record is what makes a confidence-based operating model practical. When the evidence is strong and the risk is low, an agent can be allowed to do more. When the risk is higher or the signal is unclear, the system should bring an operator back into the decision. Those responses should then improve the model over time, so autonomy expands based on performance rather than assumptions.

AUTONOMY MUST BE EARNED

Agentic AI will not remove people from payment operations. It will change where their judgment is needed. The real opportunity is to take repetitive repair work off their plates so they can focus on decisions that carry the most risk or customer impact.

That shift will not happen all at once, and that is by design. Every reviewed decision that is accepted, changed or rejected trains models and widens the band of what agents can safely do unsupervised. Banks that deploy fastest without that record get speed once. Banks that instrument the loop get compounding autonomy, where trust earned on low-risk repairs funds expansion into higher-risk ones. The moat is not the agent, but the governed track record behind it.

KAREN
STROUP

Chief Digital Officer

THE **AGENTIC ENTERPRISE**

Every conversation I've been in about agentic AI arrives at the same question: "Where do we start?" Most executives I talk to are asking which agents to deploy. The better question is what kind of company do you want to become, and how does that change everything about how work gets done?

Agentic AI is not a technology transformation. It's an enterprise transformation. The organizations that create lasting value won't be the ones with the best models. They'll be the ones willing to redesign how their companies actually operate.

From what I'm seeing, three shifts matter most.

LEAD WITH THE VISION

Don't start by asking what artificial intelligence can automate. Start by asking what your business would look like if AI were built into it from day one.

Too many organizations begin by inserting AI into existing processes. That delivers incremental efficiency, but rarely transformational value. The better starting point is to define the future state you're trying to create. If intelligent agents were part of every customer interaction, every employee experience and every major decision, how would your business operate differently? Hold the vision tightly, but the roadmap loosely.

Then work backward. Don't just optimize individual workflows — reimagine the connected work that creates value for customers. Agents don't simply automate tasks; they give us permission to redesign how work flows across the enterprise. The companies making the most progress aren't optimizing yesterday's work. They're designing tomorrow's.

BUILD THE TALENT AND CULTURE TO SCALE

Every organization needs a core of AI builders to get the flywheel moving. But they don't create transformation alone. Their role is to help business teams think differently — to question assumptions, run experiments and redesign the work they know best. The goal isn't a large, centralized AI team. It's an organization where every team has the confidence to rethink how their work gets done.

We saw this firsthand. In one of our operational functions, a small team — three people, including one who combined deep business expertise with technical fluency — began building agents against workflows they actually owned. The order mattered: Infrastructure first, then skills, then agents. That sequence, combined with the right mix of business and technical expertise, is what allowed the work to scale.

As that capability spreads, confidence spreads with it. Culture becomes the multiplier. The organizations that scale AI successfully don't just invest in technology, they build a culture where people think critically, question assumptions and are willing to experiment their way to better answers. In the agentic enterprise, the organizations that win won't be the ones that always have the right answer. They'll be the ones that learn the fastest.

REDESIGN YOUR OPERATING MODEL, NOT JUST YOUR GOVERNANCE MODEL

Agentic AI changes more than workflows, it changes how an organization operates. As agents take on more execution, people create value through judgment, creativity, relationship-building and solving complex customer problems. That requires new leadership expectations, new decision rights and new ways of measuring success.

Governance is a critical part of that operating model, but too many organizations think of it only as risk management. The other half — one we don't talk about enough — is helping people navigate constant reinvention.

Good guardrails don't just define what agents can and can't do. They create safe conditions for experimentation. They give teams clarity on where they're empowered to innovate, where human judgment is required and permission to learn when something doesn't work. The companies that move fastest won't be the ones with the most restrictive governance. They'll be the ones whose operating models make learning, adaptation and continuous reinvention part of how work gets done every day.

The destination is worth saying clearly: An enterprise where humans spend more time on the work only humans can do — judgment, creativity, the moments where context and accountability still matter. Getting there isn't a project you plan and execute. It's a muscle you build. Hold the vision tightly. Hold the roadmap loosely.

PYMNTS®

BUILDING THE
AGENT-READY PAYMENTS ENTERPRISE

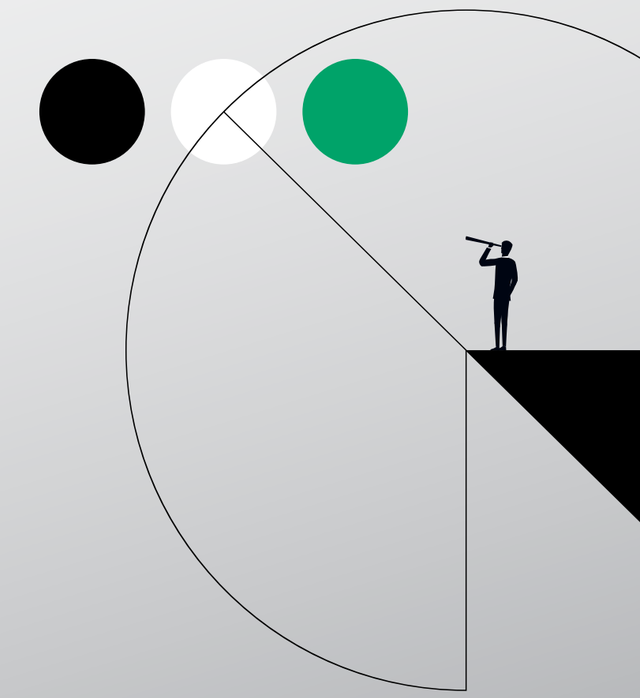

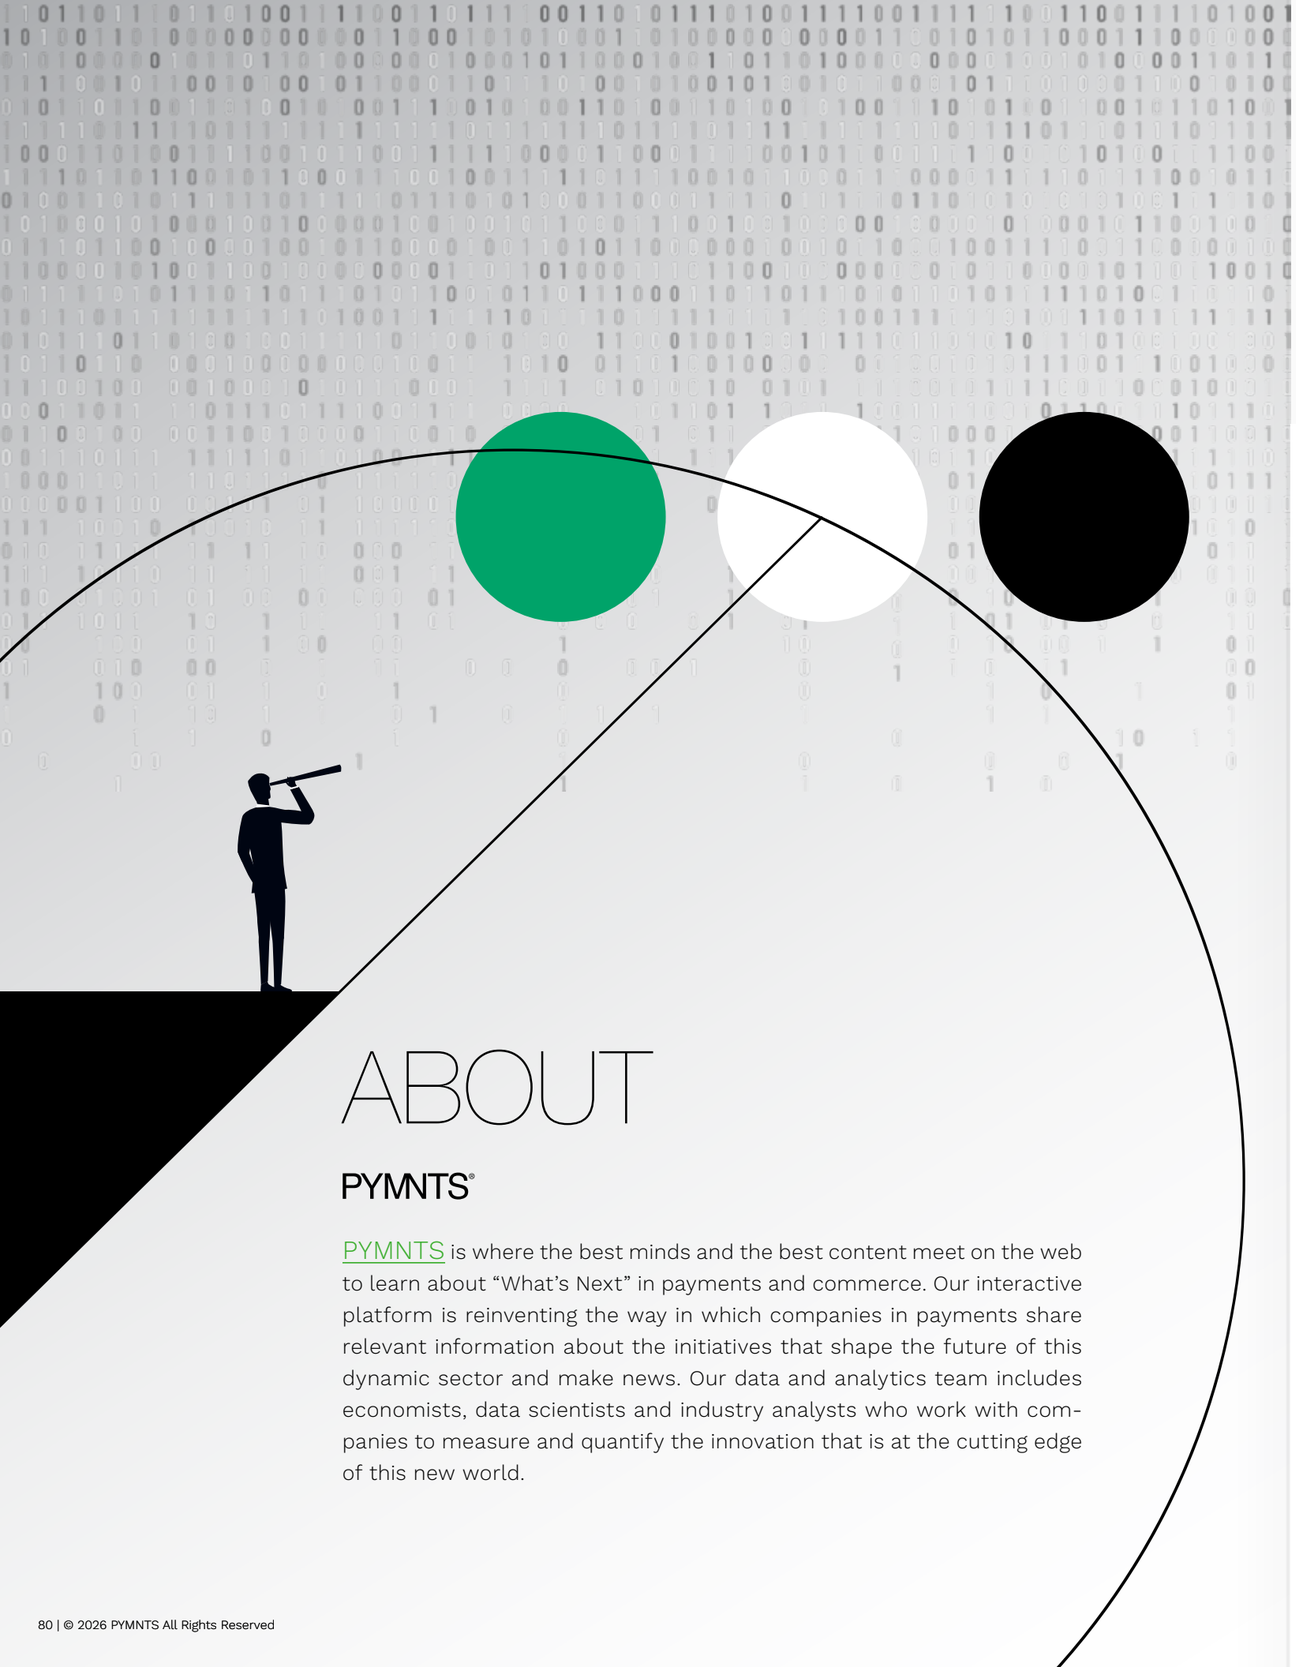

ABOUT

PYMNTS®

PYMNTS is where the best minds and the best content meet on the web to learn about “What’s Next” in payments and commerce. Our interactive platform is reinventing the way in which companies in payments share relevant information about the initiatives that shape the future of this dynamic sector and make news. Our data and analytics team includes economists, data scientists and industry analysts who work with companies to measure and quantify the innovation that is at the cutting edge of this new world.

Building the Agent-Ready Payments Enterprise eBook may be updated periodically. While reasonable efforts are made to keep the content accurate and up-to-date, PYMNTS MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND, EXPRESS OR IMPLIED, REGARDING THE CORRECTNESS, ACCURACY, COMPLETENESS, ADEQUACY, OR RELIABILITY OF OR THE USE OF OR RESULTS THAT MAY BE GENERATED FROM THE USE OF THE INFORMATION OR THAT THE CONTENT WILL SATISFY YOUR REQUIREMENTS OR EXPECTATIONS. THE CONTENT IS PROVIDED “AS IS” AND ON AN “AS AVAILABLE” BASIS. YOU EXPRESSLY AGREE THAT YOUR USE OF THE CONTENT IS AT YOUR SOLE RISK. PYMNTS.COM SHALL HAVE NO LIABILITY FOR ANY INTERRUPTIONS IN THE CONTENT THAT IS PROVIDED AND DISCLAIMS ALL WARRANTIES WITH REGARD TO THE CONTENT, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT AND TITLE. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OF CERTAIN WARRANTIES, AND, IN SUCH CASES, THE STATED EXCLUSIONS DO NOT APPLY. PYMNTS.COM RESERVES THE RIGHT AND SHOULD NOT BE LIABLE SHOULD IT EXERCISE ITS RIGHT TO MODIFY, INTERRUPT, OR DISCONTINUE THE AVAILABILITY OF THE CONTENT OR ANY COMPONENT OF IT WITH OR WITHOUT NOTICE. PYMNTS SHALL NOT BE LIABLE FOR ANY DAMAGES WHATSOEVER, AND, IN PARTICULAR, SHALL NOT BE LIABLE FOR ANY SPECIAL, INDIRECT, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, OR DAMAGES FOR LOST PROFITS, LOSS OF REVENUE, OR LOSS OF USE, ARISING OUT OF OR RELATED TO THE CONTENT, WHETHER SUCH DAMAGES ARISE IN CONTRACT, NEGLIGENCE, TORT, UNDER STATUTE, IN EQUITY, AT LAW, OR OTHERWISE, EVEN IF PYMNTS.COM HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. SOME JURISDICTIONS DO NOT ALLOW FOR THE LIMITATION OR EXCLUSION OF LIABILITY FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES, AND IN SUCH CASES SOME OF THE ABOVE LIMITATIONS DO NOT APPLY. THE ABOVE DISCLAIMERS AND LIMITATIONS ARE PROVIDED BY PYMNTS.COM AND ITS PARENTS, AFFILIATED AND RELATED COMPANIES, CONTRACTORS, AND SPONSORS, AND EACH OF ITS RESPECTIVE DIRECTORS, OFFICERS, MEMBERS, EMPLOYEES, AGENTS, CONTENT COMPONENT PROVIDERS, LICENSORS, AND ADVISERS. Components of the content original to and the compilation produced by PYMNTS is the property of PYMNTS and cannot be reproduced without its prior written permission.