

PERFICIENT

BANKING, FINANCIAL SERVICES & INSURANCE

STATE OF THE INDUSTRY

A primary-research-backed perspective on where authority over financial decisions is moving as AI scales. Built from current market evidence, Perficient practitioner interviews, and a survey of director-level and above financial services decision-makers.

01

THE MOMENT

WHERE BANKING, FINANCIAL SERVICES & INSURANCE STANDS RIGHT NOW

Three parties are already writing the rules for what machines may decide inside your institution: regulators, payments networks, and your own clients. None of them works for you. Mastercard launched Agent Pay in April 2025; Visa announced Intelligent Commerce a day later. Google's AP2 followed that September, and Visa published its Trusted Agent Protocol in October 2025. Agent-initiated payments are still early, but the operating frameworks are emerging. Real-time payment rails, including RTP and FedNow, would carry many of these transactions. These developments force institutions to answer one question: what decisions can an AI agent make, commit to, and execute without human approval?

That question is about decision authority, the boundary that determines the highest-consequence action a system may execute without human sign-off. Perficient's survey found that 76% of financial services decision-makers cannot connect AI spending to a financial outcome, and 24% have no enterprise position on where the boundary sits. Set decision authority too low, and the institution pays for AI while reserving its highest-value decisions for humans. Set it too high without evidence, and efficiency becomes financial, regulatory, and conduct exposure.

The last twelve months narrowed the framework without lightening the load. On April 17, 2026, the Federal Reserve, OCC, and FDIC issued SR 26-2, superseding SR 11-7 and SR 21-8 and placing generative and agentic AI outside model-risk guidance as novel and rapidly evolving. It is expected to matter most to Federal Reserve-regulated organizations above $30 billion in assets, and it directs banks back to their own risk-management and governance practices. Each institution now defines and defends its own controls.

Sources: *Federal Reserve SR 26-2* | *OCC Spring 2026 Semiannual Risk Perspective* | *Mastercard Agent Pay launch, April 2025* | *Visa Announces New Era of Commerce, April 2025* | *Google AP2 announcement, September 2025* | *Visa Trusted Agent Protocol, October 2025* | *FedNow/RTP Q1 2026 combined settlement data* | *Swiss Re Media Releases* | *nCino AI in Banking Report* | *Cerulli Associates Billion-Dollar RIAs Announcement* | *Cerulli Investor Skepticism Press Release* | *Berkshire Global Advisors 2026 Midyear M&A Report* | *NAIC Model Bulletin PDF* | *NAIC Catastrophe Models (Property) Hub*

PERFICIENT

SUBSEGMENT SNAPSHOTS

BANKING

Who shapes the ceiling: your examiner. SR 26-2 hands banks the discretion and keeps the accountability, so the ceiling is whatever the bank can defend at exam. Yet in a survey by banking software provider nCino, 93% of 150 U.S. banking executives cite a data-governance challenge, and only 21% tie AI investment to revenue.

PAYMENTS

Who shapes the ceiling: the networks. Agent Pay and Intelligent Commerce moved from announcement to published specification in about a year, and taking the rails means taking the spec. With no AI-specific payments regime, networks set rules you either translate into your own risk, consent, and accountability boundaries or accept as written.

WEALTH MANAGEMENT

Who shapes the ceiling: your client. No model regime governs this, so consent and fiduciary expectation set the boundary. Cerulli finds 70% of billion-dollar RIAs use AI for notetaking while 38% of affluent investors are comfortable with AI in advice. Berkshire Global Advisors counts 225 deals for RIAs with $100 million or more in assets in H1 2026, up 39%. And each one merges two consent regimes and two AI policies into a single book.

INSURANCE

Who shapes the ceiling: your regulator. The NAIC model bulletin sets the baseline, and catastrophe models are a decades-old delegation of machine authority over pricing and reserving. Secondary perils caused a record 92% of 2025's $107 billion in global insured natural catastrophe losses, per Swiss Re. Insurance already runs a high ceiling, and it holds only because carriers can explain the number to a regulator.

02

THE AI INFLECTION

THE MONEY IS COMMITTED. THE PROOF ISN'T.

Among AI use cases at pilot stage or beyond, four of the top five are back-office. Cambridge's 2026 Global AI in Financial Services Report, produced with the BIS, IMF, and World Economic Forum, found process automation (79%), data visualization (75%), and software engineering (75%) leading the list. Perficient practitioners see the same concentration: regulatory-document summarization, analyst investment research, and human-in-the-loop claims triage, running alongside back offices that still take submissions by fax. Customer support is widely adopted, while fraud, disputes, and problem resolution stay under human control.

> The next horizon of challenge will not be data scarcity, but judgment scarcity. When data is abundant, when it can be summarized, analyzed, and surfaced, then it becomes a matter of judgment. Solving the data problem is not the end goal. It is a step along the way.

BRIAN BELL

Vice President, Banking, Financial Services, and Insurance Practice, Perficient

The adoption numbers do not agree, and the disagreement is the finding. nCino puts generative AI adoption among banks at 91%. Wolters Kluwer puts production deployment at 31.8%. Cambridge puts some adoption at 81% and strategically transformative use at 14%. IDC finds financial services carries the highest AI share of DX software investment of any sector it tracks, with $151B in AI growth from 2024 to 2029, second only to retail. Institutions are investing aggressively, but most implementations remain concentrated in lower-consequence workflows. Capability is advancing faster than decision authority.

Data quality and talent led Cambridge's list of AI adoption constraints in 2020. They remain at the top in 2026, alongside privacy, explainability, human oversight, and operational resilience. Six years of investment did not move them, which is why the ceiling is set by governance rather than capability. AI capital also competes with credit and refinancing risk, fraud and cyber, compliance, and core modernization. With the federal funds target held at 3.5% to 3.75% through July 2026 and inflation above the Federal Reserve's 2% goal, every AI dollar is priced against a hurdle rate that has not come down.

Sources: nCino AI in Banking Report | Wolters Kluwer Q1 2026 Banking Compliance AI Trend Report | IDC: DX Software in Transition, AI Investment Trends by Sector | Cambridge CCAF 2026 Global

$151B

IN PROJECTED AI GROWTH FROM 2024 TO 2029

14%

REPORT STRATEGICALLY TRANSFORMATIVE AI ADOPTION

PERFICIENT'S VIEW

The institutions pulling ahead are not building better models. They're building stronger evidence behind machine-supported decisions. Evidence is what buys each tier of authority, and what defends the decision later to a regulator, a plaintiff, or a board. Most institutions already produce it and book all of it as cost.

03

THE TENSION

WHAT FINANCIAL SERVICES LEADERS TOLD US

AI remains confined to the work least likely to prove its value. Perficient's research findings reveal a self-reinforcing loop: without a clear way to price outcomes or delegate decisions, institutions default to lower-consequence use cases that provide little evidence for expanding either investment or authority.

METHODOLOGY

Perficient surveyed 50 director-level and above financial services decision-makers in August 2026 through 1Q, a mobile research platform. Findings were supplemented by insights from Perficient's Banking, Financial Services, and Insurance team.

76%

OF RESPONDENTS CANNOT CONNECT AI SPENDING TO A FINANCIAL OUTCOME.

01 AI IS BEING MEASURED BY MOTION, NOT MONEY.

An institution that cannot connect AI spend to a financial outcome cannot defend the line item in a budget cycle, and it cannot tell its board what it is buying with the delegation risk it has already accepted. Only 24% tie AI spending to revenue, margin, cost, or risk reduction. Operational and adoption metrics justify pilots. They do not justify a ceiling.

52%

OF RESPONDENTS WOULD NOT LET A MACHINE EXECUTE ANY DECISION WITHIN 24 MONTHS.

02 THE CEILING IS LOW, AND A QUARTER HAVE NOT SET ONE.

In most institutions, decision authority already exists. It simply hasn't been formally defined, with 24% of respondents reporting no enterprise position at all. Instead, the boundary gets set by whichever project team ships first, whichever vendor default goes unexamined, or whichever network specification arrives next. The institution then inherits a boundary it never approved and cannot explain.

30%

OF RESPONDENTS REJECT AI IN MATERIALLY CONSEQUENTIAL CUSTOMER OR FINANCIAL DECISIONS.

03 AI REMAINS TRAPPED IN LOW-VALUE WORK.

Caution concentrates AI in work too small to prove its own value, with only 12% of respondents prioritizing speed. The measurement problem never resolves. Withholding authority from the highest-impact use cases can feel safe, until the first supervisory exam asks who authorized the machine decisions already in production. Predefined consequence tiers give each use case a known path forward instead of a standing veto.

Source: Perficient / 1Q survey, 2026

40%

OF RESPONDENTS SLOW MODERNIZATION BECAUSE OF RISK CONCERNS OR ROUTE AROUND THE CORE.

82%

OF RESPONDENTS DIRECT THE NEXT DOLLAR TO INTERNAL CAPABILITY, FOUNDATIONS, GOVERNANCE, OR REDESIGN.

03 MODERNIZATION IS BEING DEFERRED BY CONSENSUS.

The top tier of authority needs systems that can explain and audit a decision after the fact. Institutions deferring core modernization are not postponing a technology program. They are capping how much authority they will ever be able to delegate, often without deciding to. In fact, 24% have no enterprise approach at all.

04 THE NEXT DOLLAR IS GOING INWARD.

For the next AI investment, only 6% prioritize customer-facing or revenue-generating use cases. New revenue ranks last, behind every internal priority and even behind having no agreed priority at all. Yet internal projects are the least likely to show a financial return. This internal focus also delays the safeguards customer-facing AI needs: disclosure, consent, adverse-action notices, and complaint-handling.

Source: Perficient / 1Q survey, 2026

04

THE SIGNAL

WHAT LEADERS GETTING IT RIGHT ARE DOING DIFERENTLY

The organizations making real progress are not moving faster. They sequence the work differently, and the sequence is what clears review. Across Perficient's banking, wealth, payments, and insurance work, four patterns distinguish the programs that reach production. These are practitioner observations, not survey findings.

SIGNAL 1

They make governance a design input, not a final checkpoint.

Under the old model, technology teams selected and built the tool before risk ruled on it. Review became negotiation, findings became accepted risk, and whatever shipped set the ceiling by default. Programs that scale bring legal, compliance, risk, product, and business teams together to choose the problem, tools, and deployment approach. That allows governance to take shape before there is a solution to review. The difference shows up in the pipeline: teams stop proposing use cases that will not survive.



"

In the legacy approach, technology teams would build and launch a capability first, only to have risk or compliance pull the emergency brake after the fact. High-performing organizations **bring governance into the room** before the problem is even defined.

JON PIERRE MILLAN
Director, Banking, Financial Services, and Insurance Practice, Perficient

SIGNAL 2

They fund a named decision, not a capability.

Programs that clear review fund a named process with an owner, a measured baseline, and an existing business cost. That gives finance a unit to price against compute, and it gives the board something to approve other than a platform.

"

The industry is full of **unmeasured AI excitement** right now. But when the consumption bill comes due and compute costs skyrocket, you suddenly realize you funded a capability instead of a business outcome.

JON PIERRE MILLAN
Director, Banking, Financial Services, and Insurance Practice, Perficient





SIGNAL 3

They redesign the role that signs.

When a machine drafts the recommendation, the person who signs it is doing a different job than their title describes. Institutions that have not rewritten that role are asking someone to accept accountability for a decision they no longer make. The Financial Stability Board's 2026 consultation on AI adoption warns that a human in the loop does not guarantee effective oversight and calls for clearly defined roles and responsibilities.

Source: Financial Stability Board Report

"

Leaders need to **pay more attention to the organizational structure** they want to have. They want to introduce new tools, new infrastructure, but they usually do not talk about the change in roles within their organization.

GERARDO MONTEMAYOR
Associate Vice President, Banking, Financial Services, and Insurance Practice, Perficient

SIGNAL 4

They set a velocity they can reconstruct.

Programs moving beyond pilots decide in advance which experiments enter production and which controls travel with them. The test is whether, a year later, a control function can reconstruct why a given experiment was promoted and who approved it.

"

To find **a sustainable velocity is the most important part.** A velocity where you can see the value in those AI initiatives and also adjust or adapt to it internally.

GERARDO MONTEMAYOR
Associate Vice President, Banking, Financial Services, and Insurance Practice, Perficient

05

THE QUESTION

Is anyone in your institution setting the boundary for what machines may decide, or are you inheriting it one use case at a time?

Three parties are already writing the boundaries of machine decision-making in financial services, and none of them works for you: regulators through discretion in banking and explainability in insurance, payment networks through operating specifications, and clients through consent and fiduciary expectations.

The institutions that pull ahead in 2027 will have stated in writing which decisions machines may make, who is accountable, and what evidence demonstrates better outcomes. The rest will continue operating within boundaries set by others.





Perficient is the global AI and technology consulting firm disrupting the traditional consulting model. Powered by our 7,000+ advisors, engineers, and designers, Perficient implements AI-first solutions that break conventions and deliver outcomes that matter. Proudly serving clients that represent the world's most innovative brands, and in collaboration with our powerful technology partner ecosystem, we bring deep industry expertise and data-driven design to redefine how businesses run and succeed.

Learn more at **www.perficient.com**.

DIFFERENT. FOR REAL.