
White Paper

Trust, Risk, and Reward:

Building the Identity Fabric for
Agentic AI in Financial Services

Table of Contents

A New Era of Autonomous Intelligence	3
The Evolution of AI in Financial Services	4
Identity-Enabled Autonomy: The Architecture of Agent Relationships	5
The Three Agent-Relationship Use Cases	5
IAM Capabilities that Make Autonomy Accountable	6
The Model Context Protocol: A Common Language for Trust	7
The Identity-for-AI Framework	9
Controlling LLM Model Bias	10
Industry Use Cases: From Present Reality to Agentic Futures	12
Banking	13
Insurance	16
Wealth and Asset Management	18
Building the Identity Fabric Together	20
A Partnership for Trusted Autonomy	20
From Principles to Practice: Trust, Risk and Reward	21

Trust, Risk, and Reward:

Building the Identity Fabric for Agentic AI in Financial Services

Authors:

Adam Preis, Director Product and Solution Marketing, **Ping Identity**

Deepak Goyal, Senior Manager, Cyber and Strategic Risk, **Deloitte & Touche LLP**

“Agentic AI is redefining financial services impacting client experience, product models, and efficiency ratios. Managing identity and consent across agents and humans will be fundamental to creating trust, unlocking promise of AI, and enabling growth for financial institutions.”

– **Vik Bhat, Vice Chair and Deloitte US Financial Services Leader**

A New Era of Autonomous Intelligence

The financial services industry is approaching a defining inflection point between intelligent automation and autonomous intelligence. Over the past decade, banks, insurers, and wealth managers have invested heavily in analytics, robotics, and machine learning to enhance prediction, efficiency, and control. Now, a new evolutionary stage of artificial intelligence (AI)—agentic AI—is emerging as the next step in this transformation.

This progression marks a paradigm shift from tools that merely analyse and report patterns to intelligent systems that operationalize insight, deliver customer value, reinforce compliance, and strengthen loyalty to drive sustainable profitability. More than a technological advancement, agentic AI represents a radical change in how financial institutions will plan, deliver, and scale their businesses. Within the broader industry context, this paradigm shift will challenge traditional models of financial intermediation. The opportunity is enormous. Yet with every leap in autonomy comes a corresponding rise in risk.

Agentic AI is redefining the boundaries between financial institutions, their customers, and the technology that connects them. As AI ecosystems evolve from analytic engines into autonomous actors capable of executing transactions, extending credit, or binding policies, the long-standing boundaries of accountability and oversight are being tested. Financial institutions now face a dual imperative: to realize the value of intelligent automation while ensuring that agentic AI solutions are designed, implemented, and secured in ways that preserve trust, manage risk, and meet the highest standards of regulatory and operational integrity.

This white paper explores how that imperative is reshaping financial services and sets out a joint point of view from Deloitte and Ping Identity on how institutions can design, implement, and secure the identity fabric that enables agentic AI to operate safely, responsibly, and at scale.

The Evolution of AI in Financial Services

AI in financial services has evolved through four transformative waves, each marked by broader opportunity, increased risk, and higher investment. What started as simple back-office automation has grown into advanced, AI-driven intelligence embedded across the industry.

With every wave, AI has moved closer to end customers, fundamentally changing how financial institutions operate, engage clients, and manage risks. As AI's influence broadens, so does its risk profile, making security-by-design essential in the development, deployment, and governance of modern agentic AI systems. Deloitte (see figure 1) offers a broad framework outlining this evolution through four key stages

First Wave: Rules-based Automation

Early automation depended on strict scripts to handle repetitive tasks like reconciliations and regulatory filings. While improving efficiency and scale, these tools lacked flexibility and contextual awareness, highlighting the constraints of rigid automation in a changing risk landscape.

Second Wave: Machine Learning

Machine learning brought statistical reasoning, allowing systems to spot patterns and make probabilistic predictions for uses like credit scoring and fraud detection. Although this expanded analytical capabilities, it also introduced challenges related to data quality, model bias, and governance, as human input was still needed to act on the insights.

Third Wave: Large Language Models (LLMs)

The advent of LLMs enabled AI to understand natural language and generate context-aware responses. This empowered financial institutions to deploy virtual assistants and copilots, but it also increased governance complexity due to risks such as bias, hallucinations, and potential data leakage—necessitating stricter controls.

Fourth Wave: Agentic AI

Today's agentic AI blends autonomous decision-making with reasoning. These systems can pursue objectives, select required tools, and complete complex tasks—like real-time transaction reconciliation and smart fraud detection—largely independently but within governance frameworks. This stage balances machine autonomy with institutional oversight.

Wave	Core Capability	Industry Example	Autonomy	Governance Intensity
Rules-Based Automation	Scripted workflows	Reconciliation bots	None	Low
Machine Learning	Pattern recognition	Credit-risk models	Limited	Moderate
Large-Language Models	Natural-language reasoning	Advisor copilots	Partial	High
Agentic AI	Goal-oriented planning & action	Fraud-response agents, claims bots, portfolio managers	High	Very High: requires identity, audit and human oversight

Figure 1. Agentic AI Evolution in Financial Services¹

¹ Deloitte, August 2025.

Identity-Enabled Autonomy: The Architecture of Agent Relationships

Agentic AI transforms the very notion of a user. For the first time, digital actors can make decisions, take actions, and collaborate with other systems with reduced direct human control. In this new environment, identity becomes the organizing principle of autonomy - the framework that determines who or what an agent is, what it may do, and how its actions are recorded and governed.

Ping Identity defines this new paradigm as “[Identity for AI](#)”, a foundational capability for governing the next generation of intelligent, non-human agents. Unlike traditional applications or service accounts, AI agents think, learn, and act autonomously, often finding creative ways to achieve their objectives.


They therefore form a new class of identity: intelligent, non-human entities that must be authenticated, authorized, and continuously monitored to ensure accountability and verifiable trust across every action. “Identity for AI” provides the guardrails that make this possible, embedding delegation instead of impersonation, enforcing least privilege and just-in-time access, and maintaining full traceability between the human principal, the agent acting on their behalf, and every decision the agent takes.

“AI agents are changing how business gets done.”

– Andre Durand, CEO and founder of Ping Identity.


The Three Agent-Relationship Use Cases

Every agentic AI workflow in financial services falls within one or more of three recurring relationship patterns: and on whose behalf.




Customer-to-Agent

Individuals or businesses deploy agents to interact with their financial providers. These customer-focused agents might manage payments, insurance renewals, or investment portfolios. Providers must recognise them as legitimate intermediaries, authenticate and authorize their entitlements, and apply consent frameworks that bind each agent's actions to a system of record.



Business-to-Agent

Financial institutions operate their own enterprise agents to execute internal or customer-facing tasks, from onboarding and know-your-customer (KYC) verification to fraud detection and claims processing. These agents require provisioning, supervision, and lifecycle management identical in rigor to human staff, with clear lines of accountability mapped to business owners.



Agent-to-Agent

Autonomous agents increasingly collaborate across corporate and ecosystem boundaries. A bank's fraud-detection agent may query a payments-network agent, or an insurer's claims agent may coordinate with a repair-shop agent. These interactions demand bilateral authentication, policy-based delegation, and immutable transaction records to ensure every cross-agent exchange is legitimate, limited in scope, and fully auditable.

IAM Capabilities that Make Autonomy Accountable

Across customer-to-agent, business-to-agent, and agent-to-agent use cases, a consistent set of identity and access management (IAM) capabilities creates a control plane for agentic AI. These capabilities transform identity from a static security layer into a dynamic system of governance, one that simultaneously strengthens security, enables regulatory compliance, and preserves the trust essential to financial intermediation in the era of agentic AI.

Authentication and Provenance Verification

The first step in securing autonomy is verifying who or what the agent is, and where it originates. In an agentic ecosystem, identity proofing extends beyond users and devices to include autonomous digital entities and the data they handle. Provenance verification is designed to ensure that both the agent and the datasets it consumes can be traced to legitimate, auditable sources.

For financial institutions, this means establishing cryptographic assurance of origin, verifying that a payment agent belongs to a registered customer, that an underwriting agent operates within the institution's policy domain, and that external data from partners or counterparties has not been tampered with. By anchoring every interaction to a validated identity and data source, authentication and provenance verification prevent spoofing, impersonation, and data poisoning, key threats facing future agentic AI environments.

Delegated Authorization

As agents begin to act independently, they must do so under explicit, limited mandates. Delegated authorization governs this balance by issuing short-lived, task-specific credentials that define precisely what an agent is permitted to do, under which conditions, and for how long. These credentials encode the concept of least privilege, ensuring an agent can access only the data or systems required to complete its task, nothing more.

Human-in-the-loop controls complement this model by ensuring that sensitive or high-impact actions always involve explicit human approval or post-action review. This maintains human accountability even as agents operate autonomously, reinforcing trust and transparency across all agentic activity.

In practice, this might restrict a customer's personal agent to initiating transfers under a specified threshold or confine an enterprise claims agent to processing policies in its assigned region. Delegated authorization not only reduces the attack surface but also enforces internal segregation of duties, a critical control under frameworks such as [Digital Operational Resilience Act \(DORA\)](#) and the [FCA's Operational Resilience rules](#).

Policy-Based Access Control

Policy-based access control (PBAC) applies contextual intelligence to every authorization decision, combining information about the agent's role, the sensitivity of the transaction, prevailing risk conditions, and applicable regulatory obligations. In an agentic system, PBAC replaces static access lists with real-time evaluations, allowing or denying actions based on dynamic context such as transaction value, customer vulnerability, or jurisdictional boundaries. This capability enables institutions to encode compliance policies directly into system logic, transforming regulations like the [EU AI Act](#), [Monetary Authority of Singapore \(MAS\) Pathfinder](#), or [Consumer Duty](#) into enforceable, automated controls. PBAC thus becomes the operational embodiment of both risk management and compliance assurance.

Consent Management

As AI agents act on behalf of humans, their legitimacy rests on consent. Consent management captures, stores, and enforces the permissions that link a human principal to an agent's actions. It ensures that a customer's authorization to share data or execute transactions travels securely with each request and that those permissions are revocable at any time. For financial services, this capability is central to maintaining compliance with the EU [General Data Protection Regulation \(GDPR\)](#), the [Singaporean Personal Data Protection Act \(PDPA\)](#), and similar privacy laws. It also serves as a defence against goal-hijacking or unauthorized data use, providing demonstrable proof that every agentic action aligns with the customer's explicit intent.

Lifecycle Governance

Agents, like human identities, must be created, monitored, and retired under controlled conditions. Lifecycle governance manages this process end-to-end provisioning new agents with defined roles, continuously monitoring their performance and behaviour, and de-commissioning them when they are no longer required or when risk conditions change. In regulated environments like the financial services industry, this capability provides the operational transparency supervisors expect, maintaining inventories of active agents, their access rights, and their business sponsors. It also supports internal audit and risk functions by ensuring that dormant or orphaned agents cannot persist unseen within critical systems.

Auditability and Non-Repudiation

Finally, every autonomous action must be observable and reconstructable. Auditability and non-repudiation maintain an immutable, time-stamped record of all agent activities, linking each event to the credentials, consent tokens, and policy decisions that authorized it. These records enable institutions to satisfy regulatory demands for explainability under the EU AI Act and for incident traceability under DORA and MAS Pathfinder. More importantly, they provide the evidence base that sustains trust: customers, regulators, and counterparties can see not only what an agent did, but also that it acted under legitimate authority and within approved boundaries.

The Model Context Protocol: A Common Language for Trust

As the number of autonomous agents and connected systems multiplies across financial institutions, managing identity, access, and trust through proprietary or point-to-point integrations becomes unsustainable. Each agent-to-agent or cross-domain interaction, whether between a customer's personal assistant, a bank's fraud-detection engine, or a third-party settlement network, requires a consistent method to authenticate participants, exchange credentials securely, and enforce policy boundaries.

The [Model Context Protocol](#) (MCP) has emerged to meet this need. It provides a standards framework that links the core IAM capabilities of authentication, authorization, consent, and audit across distributed, multi-agent environments. By extending established security protocols such as OAuth 2.0 and OpenID Connect, MCP establishes a common language for trust, enabling secure discovery, credential exchange, and contextual communication among agents, tools, and data sources. Within financial-services ecosystems, it allows a customer's personal agent, a bank's operational agent, and a third-party service agent to recognise and validate one another, enforcing mutual trust and explicit scopes of authority.

In this model, identity becomes an active participant in every transaction rather than a one-time verification step. Authentication confirms origin; authorization defines scope; audit preserves evidence. Through this triad, institutions can ensure that each autonomous action, whether initiated by a human, a corporate system, or an AI agent, remains within approved boundaries and traceable to an accountable entity.




Figure 3. MCP as the common language for trust.

The Identity-for-AI Framework

As agentic AI becomes more pervasive, standards like the MCP provide a technical foundation for interoperability and trust. Yet standards alone are not sufficient; financial service providers also need a design philosophy that defines how identity safeguards should be applied to autonomous actors.

Ping Identity seeks to provide that very foundation through its '**Identity for AI**' solution. It translates abstract governance concepts into actionable design principles, ensuring that autonomy in financial systems operates under continuous human accountability. This solution rests on five guiding tenets that together enable safe, compliant, and scalable deployment of agentic AI across banking, insurance, and wealth management.

Tenet 1

Delegation Over Impersonation

In financial services, accountability depends on knowing who performed an action and under whose authority. Agents therefore act on behalf of humans but never as them. Each agent, whether a personal banking assistant, a digital claims adjuster, or a trading algorithm, receives its own cryptographically verifiable identity and permission set, distinct from the individual or enterprise it represents.

Delegation ensures that when an agent initiates a payment, approves a claim, or executes a trade, its credentials trace back to an accountable principal but cannot be used to impersonate that person or institution. This separation of identities eliminates credential sharing, a leading cause of insider and supply-chain breaches and forms the cornerstone of digital trust.

Tenet 2

Least Privilege & *Just-in-Time* Access

Financial institutions operate under stringent principles of compliance oversight: the same must apply to agentic AI. Agents are granted only the minimum access required to perform a specific task, and only for the precise duration necessary to complete it. This principle of least privilege limits lateral movement and data exposure if an agent is compromised, while just-in-time (JIT) access reduces persistent risk by removing standing credentials altogether.

In practice, this means a loan-underwriting agent can view but not export customer data, a claims-processing agent can approve reimbursements only within its monetary threshold, and an advisory agent's access to client portfolios expires immediately after execution. By bounding privilege in time and scope, institutions reduce systemic exposure and strengthen compliance with regulatory expectations for risk containment.

Tenet 3

Human-in-the-Loop

Even in the most advanced agentic systems, humans remain the final arbiters of fiduciary and ethical accountability. High-impact, high-risk, or anomalous agentic actions, such as large fund transfers, unusual insurance payouts, or cross-border data access, must trigger human-in-the-loop confirmation before completion.

This safeguard mirrors the operational controls long familiar in financial supervision: dual authorisation, four-eyes review, and escalation on exception. Embedding these triggers directly into agentic workflows ensures that oversight occurs at the speed of automation, not after the fact.

Tenet 4

Continuous Monitoring & Traceability

Autonomous systems require continuous visibility to remain trustworthy. Every decision, data request, and action executed by an agent must be logged with complete contextual metadata, who authorized it, what data was used, what outcome resulted, and which controls were applied. Continuous monitoring transforms this data into intelligence: it allows institutions to detect anomalous behaviour, validate compliance, and feed learnings back into model risk management.

For regulators, this traceability delivers the transparency they increasingly demand; for customers, it provides the reassurance that even autonomous operations remain observable and accountable. In effect, monitoring and traceability are the audit trail of the agentic era.

Tenet 5

Lifecycle Governance

Agents, like human employees, require full lifecycle oversight, from creation and onboarding through supervision, modification, and eventual decommissioning. Lifecycle governance ensures that each agent's purpose, permissions, and accountability chain are defined before deployment, continuously reviewed throughout its operation, and formally retired when its function or context changes.

This capability addresses a critical control gap in emerging AI regulation: the requirement to maintain a complete and current inventory of operational AI systems. By extending identity governance to non-human identity (NHI), financial institutions can meet supervisory expectations for model inventories, access reviews, and policy attestation while preventing “ghost agents” from persisting unseen in production environments.

Applied consistently, the ‘Identity for AI’ solution transforms IAM from a reactive compliance function into the enabling architecture of trusted autonomy. It allows financial institutions to innovate at speed, confident that every autonomous decision is authenticated, authorised, observed, and attributable, a complete chain of trust from design through execution.

Agentic AI and the Compliance Convergence

As financial institutions embed autonomous decisioning, regulators worldwide are aligning on one key concept: autonomy must not dilute accountability. New frameworks, from the EU AI Act and DORA to the UK's FCA's operational-resilience rules and MAS Pathfinder in Singapore, demand that governance be built into the systems themselves. Compliance is shifting from retrospective reporting to architectural assurance.

Identity serves as a common denominator, helping financial institutions achieve compliance as regulatory expectations accelerate and diversify. Authentication, authorization, consent, and audit will become the shared language through which transparency, oversight, and explainability in the agentic AI era are proven. Regardless of geography or regulation, every AI agent must be identifiable, its actions traceable, and its permissions bounded in time and scope.

This identity fabric makes compliance continuous and self-evidencing. Provenance verification validates data integrity; delegated authorization limits autonomy; policy-based access control embeds regulation into code; and lifecycle governance maintains visibility and auditability. Regulation becomes a living control plane rather than a static checklist.

By embedding these identity safeguards, financial institutions turn compliance from a constraint into an enabler. What Deloitte calls "compliance by design" replaces friction with confidence, allowing agentic AI to scale safely, demonstrate accountability in real time, and preserve the trust that underpins financial stability.

Controlling LLM Model Bias

Agentic AI, including LLMs, enables autonomous decision-making in financial services, powering product recommendations, client advisory, compliance checks, and risk assessments. However, LLMs are prone to bias, creating significant risk in areas of compliance, reputation, and financial stability.

Why Model Bias Is a Serious Concern




Trust & Compliance

Unbiased outputs are essential for regulatory compliance and maintaining client trust.



Reputational Risk

Biased decisions can lead to negative media, client impact, and legal ramifications.



Systemic Risk

Bias at scale can affect financial markets and perpetuate inequalities.

To address these challenges, a multi-pronged approach is recommended. Rigorous bias assessment should begin at the model design stage, employing diverse and representative training datasets and utilizing synthetic data to fill coverage gaps, particularly in sensitive or underrepresented categories. Regular auditing of workflow prompts helps to mitigate inadvertent bias in practical deployments.

Advanced monitoring and testing are essential, including implementing domain-specific bias detection harnesses and periodic third-party audits to ensure ongoing fairness. Human-in-the-loop controls remain critical, particularly for high-impact decisions like lending or customer segmentation; expert review and override protocols act as safeguards against automated errors. Furthermore, embedding explainability and transparency, through explainable AI techniques and comprehensive documentation, empowers both stakeholders and regulators to understand, audit, and challenge AI-driven outcomes.


INDUSTRY USE CASES:

From Present Reality to Agentic Futures

The financial-services industry has reached the point where efficiency alone no longer equates to progress. Years of process automation and digitization have delivered scale, but they have also multiplied exposure, expanding the attack surface, fragmenting oversight, and straining regulatory control. Malicious actors now exploit the same digital channels used by customers, while data-privacy obligations and operational-resilience rules grow tighter with every incident. As banks, insurers, and wealth managers begin to deploy agentic AI, the central question is no longer whether to automate decisions, but how to do so securely, efficiently, and effectively.

The answer lies in establishing an **identity fabric**: a cohesive layer that unites security, trust, and compliance across human and agentic interactions. This fabric integrates best-of-breed authentication, authorization, consent, policy enforcement, and audit into every interaction, ensuring that each autonomous action can be verified, constrained, and traced. It transforms identity from a defensive safeguard into an active governance framework: preventing misuse before it occurs, proving compliance as it happens, and enabling innovation that regulators and customers alike can trust.

By embedding security-by-design principles into the core of agentic AI, the identity fabric allows financial institutions to contain risk while unlocking the productivity, personalization, and customer-experience gains of autonomous operations. The following use cases illustrate how these capabilities can deliver measurable value.



Autonomous Fraud Detection and Response



Challenge

Fraud remains one of the world's largest operational risks, while authorized push payment (APP) fraud alone is projected to reach over \$8.3 billion² globally in 2024. Static rules and manual triage are slow, reactive, and prone to human error. Customers are losing confidence, regulators continue to demand faster restitution, and operational teams face alert fatigue, as compliance requirements increase.

Target State

Agentic AI fraud defense establishes a continuous, identity-anchored control loop. Authenticated business-side fraud agents monitor payments in real time and, when anomalies arise, initiate bilateral checks with customer-agents. Through delegated authorization and customer-driven consent, both sides validate intent before execution. Transactions proceed or pause automatically, and every event is captured immutably for audit and restitution.

Value Drivers

Agentic fraud management reframes compliance and loss prevention as a single, continuous process of protection and assurance. Deloitte research on intelligent automation shows that banks applying advanced AI to operational risk may be able to materially lower fraud losses and associated handling costs by compressing investigation time and improving accuracy. To realize value, financial service decision-makers should:

- ☐ Define a unified fraud-agent architecture tied to their identity fabric.
- ☐ Establish customer-agent channels for instant transaction confirmation.
- ☐ Integrate delegated-authorization limits aligned to risk thresholds.
- ☐ Automate evidence capture for compliance reporting.

² Deloitte, October 2025.

Identity for AI Capabilities

- ✓ **Authentication and Provenance Verification:**
Confirms the legitimacy of each agent and data source across payment networks.
- ✓ **Delegated Authorization:**
Limits transaction authority by amount, time, and channel.
- ✓ **Policy-Based Access Control (PBAC):**
Applies contextual risk scoring and dual approval for high-risk cases.
- ✓ **Consent Management:**
Records customer confirmation as enforceable consent artifacts.
- ✓ **Auditability and Non-Repudiation:**
Produces continuous evidence of every decision and action.

Human-in-the-Loop Controls

Even in real-time agentic fraud prevention, human oversight remains essential to maintaining trust, accountability, and regulatory assurance. Fraud agents can identify and act on anomalous behaviour instantly, but final adjudication on disputed or high-value cases needs to rest with trained analysts operating under defined escalation policies. Human-in-the-loop controls provide that safeguard, ensuring that exceptions, customer disputes, or policy ambiguities are reviewed and resolved by accountable personnel before funds are released or decisions finalized.

Personalised Financial-Advisory Agents



Challenge

Advisory models are constrained by manual data gathering and episodic engagement. Advisors spend more time compiling data than providing guidance. Privacy laws demand explicit consent and full traceability, requirements that make continuous personalisation difficult without new controls. To add to this, many customers feel uncomfortable with giving human and agentic AI advisors unchecked access to their accounts and data.

Target State

A network of intelligent, identity-anchored agents enables continuous, compliant personalization. The customer-agent securely gathers financial goals and behavioural insights under explicit consent, while the provider-advisory agent interprets this information through the institution's policy and risk engines, applying embedded suitability and disclosure rules. Each recommendation is traceable from data source to delivery and validated by a human advisor, preserving fiduciary oversight while delivering always-on, transparent advice that strengthens both trust and client engagement.

Value Drivers

Embedding identity control enables hyper-personalization that is demonstrably compliant and secure. To realize value, financial service decision-makers should:

- ☐ Use consent tokens as compliance artifacts within CRM and advisory systems.
- ☐ Automate suitability checks via PBAC tied to investment-policy rules.
- ☐ Provide advisors with real-time transparency dashboards for agent activity.
- ☐ Capture and store advisory communications for compliance audits.
- ☐ Measure impact by client-engagement frequency and regulatory-review efficiency.

Identity for AI Capabilities

- ☒ **Consent Management:**
Captures and enforces data-sharing permissions for each customer.
- ☒ **Policy-Based Access Control:**
Applies suitability and disclosure rules in real time.
- ☒ **Delegated Authorization:**
Separates analysis rights (agent) from execution rights (advisor).
- ☒ **Continuous Monitoring and Traceability:**
Logs every recommendation and approval event.
- ☒ **Lifecycle Governance:**
Reviews and retires advisory agents as client portfolios evolve.

Human-in-the-Loop Controls

Human oversight remains indispensable in agentic advisory models, where decisions touch customers' financial wellbeing and regulatory obligations. While advisory agents can analyse data and generate recommendations autonomously, final approval must come from licensed advisors who apply professional judgment and contextual understanding before execution. These human-in-the-loop controls preserve fiduciary accountability, ensuring that every automated insight is interpreted through the lens of client suitability and regulatory compliance. They also reinforce customer trust.

Agentic Loan Underwriting and Credit Scoring



Challenge

Credit underwriting remains one of the most fragmented and risk-exposed workflows in banking. Manual document verification, repeated data entry, and opaque decision logic prolong approval times and increase operational cost. Model risk governance is under pressure as regulators, particularly under the EU AI Act, demand transparency, explainability, and human accountability for automated credit decisions. The result is a process that is expensive to operate, vulnerable to bias, and slow to deliver outcomes customers now expect in real time.

Target State

Leverage an end-to-end network of authenticated agents coordinating underwriting decisions. In this use case, a customer-agent compiles verified digital credentials such as identity, income statements, and open-banking data, all authenticated through provenance verification. The bank's underwriting agent then connects with credit-bureau agents via the MCP, using delegated rights to gather only the information authorized for that specific application. PBAC helps ensure decisions adhere to jurisdictional lending rules and internal risk limits, while every data exchange and decision is immutably recorded for audit.

Value Drivers

Agentic underwriting delivers transparent, explainable credit decisions at digital speed, turning a legacy compliance obligation into a competitive advantage. To realize value, financial service decision-makers should:

- ☐ Adopt verifiable-credential standards for borrower data exchange.
- ☐ Establish MCP-based agent communications with credit bureaus and fintech data providers.
- ☐ Automate explainability reporting aligned to local compliance requirements.
- ☐ Apply delegated-access tokens that expire at decision issuance to prevent data over-exposure.
- ☐ Quantify return through higher approval accuracy, faster turnaround, and reduced audit effort.


Identity for AI Capabilities

- ☒ **Authentication and Provenance Verification:**
Confirms borrower identity and ensures that all external data originates from validated, auditable sources.
- ☒ **Delegated Authorization:**
Limits the underwriting agent's access and data retrieval strictly to the duration and context of the loan application.
- ☒ **Policy-Based Access Control (PBAC):**
Applies jurisdictional and product-specific lending policies dynamically during decisioning.
- ☒ **Auditability and Explainability:**
Maintains complete, immutable records of decisions and the reasoning behind them for regulator and auditor review.
- ☒ **Lifecycle Governance:**
Provides continuous oversight of underwriting agents, including provisioning, monitoring, and de-commissioning as models or policies evolve.

Human-in-the-Loop Controls

Even in real-time agentic fraud prevention, human oversight remains essential to maintaining trust, accountability, and regulatory assurance. Fraud agents can identify and act on anomalous behaviour instantly, but final adjudication on disputed or high-value cases needs to rest with trained analysts operating under defined escalation policies. Human-in-the-loop controls provide that safeguard, ensuring that exceptions, customer disputes, or policy ambiguities are reviewed and resolved by accountable personnel before funds are released or decisions finalized.

Automated Claims Management



Challenge

Claims operations remain one of the most expensive and customer-visible functions in insurance. Manual review, fragmented workflows, and inconsistent fraud checks cause long cycle times, rising loss-adjustment expense, and high regulatory exposure. Regulators expect faster resolution, transparent reasoning, and human oversight, while customers increasingly judge insurers by how quickly and fairly, they settle claims. Legacy systems make it difficult to meet these demands without increasing operational risk.

Target State

Agentic automation underpinned by the identity fabric connects every step of the claims lifecycle in a single, governed workflow. The customer-agent gathers verified evidence and transmits it to the insurer's claims-provider-agent using authenticated credentials. The provider-agent validates policy coverage through provenance verification and engages external partner agents, for example, medical networks or repair shops, through delegated authorization. PBAC enforces payment thresholds and fraud checks automatically, while lifecycle governance ensures active claims agents remain supervised and retired when no longer required. All actions are logged through auditability and non-repudiation, creating immutable proof of compliance and fairness.

Value Drivers

Identity-anchored automation cuts cost and risk simultaneously while making compliance demonstrable. Leading insurance providers are increasing their investments in AI to shorten average settlement, including through the use of their insurance copilots. To realize value, financial service decision-makers should:

- ☐ Deploy end-to-end provenance verification for all claim submissions.
- ☐ Authorize automated payouts within defined thresholds using delegated tokens.
- ☐ Provide compliance teams with live audit dashboards for claim status and exceptions.
- ☐ Extend PBAC and consent management across all insurance policies.
- ☐ Measure success through reduction in fraud losses and regulatory-response times.


Identity for AI Capabilities

- ☒ **Authentication and Provenance Verification:**
Confirms identity of claimants and third-party partners, protecting against fraudulent submissions.
- ☒ **Delegated Authorization:**
Restricts payout authority by amount, product type, and jurisdiction.
- ☒ **Policy-Based Access Control (PBAC):**
Enforces claim-handling and anti-fraud policies in real time.
- ☒ **Lifecycle Governance:**
Tracks active claims agents and their entitlements across lines of business.
- ☒ **Auditability & Non-Repudiation:**
Captures every action for compliance traceability requirements.

Human-in-the-Loop Controls

In claims management, automation can accelerate resolution, but human oversight remains essential to ensuring fairness, accuracy, and compliance. While agentic systems can validate evidence, calculate payouts, and flag anomalies, high-value or disputed claims must be escalated for human review to preserve impartial judgement and customer confidence. Human-in-the-loop controls allow claims specialists to assess complex cases, override or confirm agentic recommendations, and document rationale for regulators and auditors.

Dynamic Risk Pricing and Policy Recommendations



Challenge

Pricing models are largely static, recalibrated annually and often using outdated data. This lag exposes insurers to underwriting risk and limits competitiveness. Privacy regulations restrict the use of telemetry data, and regulators require transparency and fairness in premium determination, creating a compliance bottleneck for real-time pricing innovation.

Target State

A governed network of customer-agents and provider-pricing agents delivers continuous, explainable pricing. Customer-agents transmit telemetry or behavioural data under explicit, revocable consent. Pricing agents apply PBAC and delegated authorization to use only permitted data for exposure calculations. Every rate change is validated through auditability, logged for supervision, and shared securely with reinsurer or distribution-partner agents via MCP connections.

Value Drivers

Continuous pricing supported by identity governance can deliver fairer outcomes for customers and sustainable margins for insurers. To realize value, financial service decision-makers should:

- ☐ Implement consent-management frameworks aligned with local privacy requirements.
- ☐ Automate audit reporting for rate adjustments and customer notifications.
- ☐ Use delegated credentials for premium updates under controlled thresholds.
- ☐ Integrate reinsurer and broker agents for transparent exposure visibility.
- ☐ Monitor fairness metrics to evidence compliance and customer trust.

Identity for AI Capabilities

- ☒ **Consent Management:**
Captures ongoing permission for telemetry and behavioural data use.
- ☒ **Policy-Based Access Control:**
Applies actuarial and compliance constraints dynamically.
- ☒ **Delegated Authorization:**
Limits rate adjustments to authorized percentages or timeframes.
- ☒ **Continuous Monitoring and Traceability:**
Generates instant audit trails for supervisory review.
- ☒ **Lifecycle Governance:**
Ensures pricing agents remain certified and compliant with solvency standards.

Human-in-the-Loop Controls

Dynamic pricing introduces continuous recalibration of premiums, making human oversight vital to maintaining fairness, transparency, and regulatory confidence. While pricing agents can analyse telemetry and behavioural data in real time, human actuaries and risk managers must review algorithmic adjustments to ensure that rates remain justifiable, proportionate, and non-discriminatory. Human-in-the-loop controls provide this safeguard, verifying that agentic decisions align with underwriting principles, consumer-protection standards, and solvency requirements.

Autonomous Portfolio Rebalancing



Challenge

Manual rebalancing exposes portfolios to drift and delays, while the need to document every decision for MiFID II and other regulatory requirements strains operational resources. Advisors often spend time reconciling records rather than managing strategy.

Target State

The identity fabric enables a provider-to-agent and customer-to-agent configuration for continuous alignment. Portfolio management agents monitor exposures in real time and execute small trades autonomously under delegated authorisation, escalating significant reallocations through human-in-the-loop approval. PBAC applies risk and suitability parameters to every trade, and auditability automatically records rationale and outcome for supervisory review.

Value Drivers

Continuous, compliant rebalancing enhances both performance and governance, reducing tracking errors and operational costs. To realize value, financial service decision-makers should:

- ☐ Map investment-policy rules directly into PBAC logic.
- ☐ Configure delegated-access tokens for automatic rebalancing under mandate.
- ☐ Automate MiFID II and SEC reporting from audit logs.
- ☐ Provide advisors with dashboards for oversight and approvals.
- ☐ Track compliance efficiency via reduced review time and error rates.

Identity for AI Capabilities

- ✓ **Delegated Authorization:**
Defines trade thresholds and escalation requirements.
- ✓ **Policy-Based Access Control:**
Embeds risk and suitability rules directly in the trading workflow.
- ✓ **Human in the Loop:**
Ensures manual validation for high-impact trades.
- ✓ **Auditability and Explainability:**
Creates traceable evidence for MiFID II and other regulatory reviews.
- ✓ **Lifecycle Governance:**
Oversees portfolio-agent roles, entitlements, and retirements.

Human-in-the-Loop Controls

In portfolio management, human oversight is indispensable for maintaining fiduciary responsibility and regulatory assurance. While portfolio management agents can monitor positions and execute small adjustments autonomously, larger reallocations that materially affect a client's risk profile or investment objectives require human validation. Advisors act as the critical checkpoint, reviewing rationale, confirming suitability, and approving execution before orders are placed.

Wealth Concierge Agents



Challenge

High-net-worth clients expect 24/7 personalized service, yet heightened privacy obligations make responsiveness risky without granular control. Without verifiable identity boundaries, client-facing automation risks breaching confidentiality or mis-executing transactions.

Target State

A combination of customer-agents and concierge-provider-agents enables always-available, identity-secured service. Customer-agents initiate requests through adaptive authentication; concierge agents fulfil them under policy-based access control, escalating sensitive actions through human-in-the-loop review. Consent management binds each request to an approved principal, while lifecycle governance ensures privileges remain appropriate to client tier and relationship status.

Value and Decision Points

Identity-anchored concierge automation expands capacity while strengthening trust. RBC Wealth's pilot reported significant advisor-productivity gains and client-satisfaction increases arising from its agentic transformation initiatives.³ To realize value, financial service decision-makers should:

- ☐ Deploy adaptive-authentication and behavioural analytics for client verification.
- ☐ Use PBAC to enforce advisory and transaction segregation.
- ☐ Capture consent tokens for all high-risk actions.
- ☐ Automate compliance evidence extraction from audit logs.
- ☐ Quantify advisor-capacity uplift and compliance-audit savings.

³ Salesforce

Identity for AI Capabilities

- ✓ **Adaptive Authentication:**
Verifies clients via behavioural and contextual analysis.
- ✓ **Policy-Based Access Control:**
Segregates information from transactional functions.
- ✓ **Consent Management:**
Confirms approval for advisory or monetary actions.
- ✓ **Lifecycle Governance:**
Adjusts agent permissions as client tiers change.
- ✓ **Auditability:**
Preserves end-to-end records for privacy and conduct compliance.

Human-in-the-Loop Controls

Even in the most advanced concierge models, human oversight is essential to preserve client trust, privacy, and regulatory compliance. While concierge agents can securely fulfil informational and routine service requests, any action involving advice, fund movement, or changes to client permissions must trigger human review. Advisors serve as the final layer of validation, confirming client intent, verifying regulatory suitability, and ensuring that sensitive instructions are executed correctly.

Building the Identity Fabric Together

Agentic AI is changing what it means to manage risk and deliver value in financial services. Institutions can no longer separate innovation from governance: the same systems that enable autonomy must also prove accountability. This convergence of capability and control is the foundation of the identity fabric, the shared architecture of security, compliance, and trust that underpins the safe adoption of agentic AI. It is within this space that Deloitte and Ping Identity combine their strengths to help the industry move from experimentation to enterprise-scale deployment.

“With Identity for AI, we give organizations the guardrails to innovate responsibly and with confidence through enterprise-grade identity management. Identity is becoming the universal language of accountability—for humans and agents alike. It’s how every AI decision can be verified and trusted.”

– **Andre Durand**, CEO and founder of Ping Identity.

Collaborating to Create Trusted Autonomy

Together, Deloitte and Ping Identity assist financial institutions to operationalize the Identity for AI framework through design, policy, and technology. Deloitte brings strategic and regulatory experience: defining target operating models, mapping regulatory obligations, and helping clients align governance with enterprise risk appetite.

Ping Identity provides the technical foundation through the Ping Identity Platform unifying authentication, authorization, consent, policy enforcement, and audit across customer, workforce, third-party, and non-human identity use cases. This alliance aims to help clients at every stage of adoption, from initial proof of concept to scaled production, advance within a verifiable framework of compliance and assurance.

From Principles to Practice: Trust, Risk and Reward

The business case for identity-enabled autonomy is founded on three inseparable outcomes: trust, risk, and reward, each reinforced by the capabilities described throughout this white paper. Together, they form the strategic logic for building the identity fabric that allows financial institutions to securely embrace agentic AI at scale.

Trust

In a digital economy where autonomy challenges accountability, trust becomes the defining currency. The identity fabric anchors every agentic interaction in verified identity, explicit consent, and continuous traceability, ensuring that each decision, transaction, and data exchange is authentic, authorized, and auditable. This transparency rebuilds confidence among customers, regulators, and counterparties, transforming compliance from obligation into assurance.

Risk

As agentic AI scales in the industry, the ability to govern agents in real time determines resilience. Embedded controls, authentication, delegated authorization, policy enforcement, and lifecycle governance, contain systemic exposure, detect anomalies early, and enforce regulatory guardrails automatically. Risk management shifts from retrospective correction to continuous prevention, aligning security and compliance in a single operational framework.

Reward

When trust and risk are managed through identity, innovation scales safely. Agentic AI governed by the identity fabric accelerates product development, reduces cost-to-serve, and improves speed to market while maintaining verifiable oversight. The result is a sustainable growth engine: operational efficiency that compounds into competitive differentiation, brand trust, and enduring stakeholder value.

Taken together, **trust, risk, and reward** define the new equilibrium for financial services in the age of agentic AI. Institutions that embed these principles into their identity fabric move beyond controlled experimentation to scalable, repeatable, and measurable transformation, where autonomy is not merely permitted but trusted, governed, and profitable.

Continue your agentic AI journey with Ping Identity and Deloitte⁴ today by getting your free [Ultimate Guide for Identity for AI](#).

⁴As used in this document, "Deloitte" means Deloitte & Touche LLP, a subsidiary of Deloitte LLP. Please see www.deloitte.com/us/about for a detailed description of our legal structure. Certain services may not be available to attest clients under the rules and regulations of public accounting.