WHITE PAPER

AI-Driven Financial Claims Automation with Pega

How governed AI agents and workflow automation transform insurance claims and banking disputes



Varun Vemulapalli

November 2026

Executive Summary

AI can make financial claims faster, cheaper and more consistent, but only when it runs inside a governed workflow that enforces regulatory deadlines and keeps people accountable for consequential decisions. Pega provides that workflow layer for both insurance claims and banking disputes.

- **The problem.** Claims combine unstructured evidence, fraud risk and fixed legal deadlines. Federal Reserve examiners ranked Regulation E error-resolution failures among the top violations they found in 2024.
- **The opportunity.** Controlled studies show productivity gains of 14% to 40% or more on suitable tasks. McKinsey estimates that more than half of current claims activities could be automated by 2030.
- **The platform.** Pega combines case management, rules, Process AI, the Blueprint design tool and, since Infinity 26 (July 2026), governed AI agents with MCP connectivity and per-case pricing. Packaged applications such as Smart Dispute and Pega Claims come with network and regulatory rules built in.
- **The risk.** Gartner expects more than 40% of agentic AI projects to be cancelled by the end of 2027. Over-reliance, prompt injection, unfair outcomes and vendor failures are the main hazards.

Recommendations: start with one high-volume claim type; design it in Blueprint on a packaged Pega application; keep deadlines and approvals in the case rather than in AI prompts; run in shadow mode before granting autonomy; govern under the NAIC bulletin, Regulation E and the NIST AI RMF; and measure cycle time, cost per claim, regulatory exceptions and fairness from day one.

1. Introduction

A claim is the moment when a financial institution proves whether its promise to a customer is worth anything. For an insurer, that promise is payment after a loss. For a bank or card issuer, it is protection when a transaction is wrong, unauthorized or fraudulent. In both cases the customer is often stressed, the money at stake is real, and regulators set the rules and deadlines.

This paper uses the term *financial claims* to cover two related families of work: insurance claims (property and casualty, life, and disability) and banking claims (card disputes, chargebacks, fraud claims and payment investigations). They share a common shape: intake of a customer's assertion, evidence gathering, a decision against policy or regulation, a financial adjustment, and correspondence, all under a deadline.

The paper examines how artificial intelligence, orchestrated by a workflow platform, can make that process faster, more accurate and more defensible. It uses Pega (Pegasystems) as the reference automation platform because Pega combines case management, rules, AI decisioning and agent

orchestration in one product line, and offers packaged applications for both insurance claims and banking disputes. The analysis and recommendations are the author's own; product capabilities are described from Pega's published materials and independent coverage, and should be confirmed against the version an organization licenses.

2. The Financial Claims Challenge

Financial claims are hard to automate because they combine unstructured evidence, adversarial behavior and fixed legal deadlines in a single process. The main claim families differ in detail but share these pressures.

Claim family	Typical trigger	Main rulebooks	Key pressure
Property and casualty insurance	Auto accident, home damage, liability	Policy wording, state insurance law, unfair claims practices rules	Large volumes after catastrophes; leakage and fraud
Life and disability insurance	Death, illness, inability to work	Policy terms, state law, beneficiary rules	Sensitive customers; years of policy history to research
Card disputes and chargebacks	Unrecognized charge, goods not received	Visa, Mastercard and Amex network rules; Reg Z (credit)	Network rules change often; strict response windows
Debit and electronic payment errors	Unauthorized transfer, wrong amount	Regulation E (US); equivalents such as Section 75 in the UK	Statutory investigation and provisional credit deadlines
Payment investigations and exceptions	Failed, delayed or misrouted payment	Scheme and correspondent bank rules	Manual research across many systems

2.1 Fraud

Fraud raises the stakes on every claim decision. The Coalition Against Insurance Fraud estimated in 2022 that insurance fraud costs the United States $308.6 billion a year, with life insurance the largest category at $74.7 billion (Coalition Against Insurance Fraud, 2022). That figure is contested: a 2024 paper from the Rutgers Center for Risk and Responsibility argues the estimate cannot be defended (Center for Risk and Responsibility, 2024). Whatever the exact size, automation must speed up honest claims without making fraud easier, and must not treat honest claimants as suspects.

2.2 Regulatory clocks

Banking claims run on statutory timers. Under Regulation E, a US institution generally has 10 business days after a notice of error to investigate. It may extend to 45 calendar days only if it provisionally credits the consumer within those 10 business days, and some cases allow 20 or 90 days (Benton, 2025). The institution must report results within three business days of finishing the investigation (Electronic Code of Federal Regulations, n.d.).

These rules are routinely missed. Federal Reserve examiners listed Regulation E error-resolution failures, including late provisional credit and late notices, among the top violations found in 2024. They traced several of these failures to weak procedures and controls at banks and at their third-party vendors (Federal Reserve System, 2026). This is a workflow problem as much as a legal one: deadlines must be computed, tracked and enforced on every case.

3. The Evolution of Claims Automation

Claims automation has moved through four generations, and a modern claims operation on Pega typically uses all four in the same case, choosing the simplest technique that reliably handles each step.

Generation	What it does in claims	How it appears in Pega	Limitation
1. Rules and BPM	Routes work, computes deadlines, applies coverage and network rules	Case types, stages, service-level agreements, business rules	Cannot read unstructured evidence or weigh judgment calls
2. Robotic process automation	Keys data into legacy claim and core banking screens	Robotic automation used alongside cases	Brittle when screens change
3. Predictive AI	Scores complexity, fraud risk and likely outcomes	Process AI and adaptive models inside the workflow	Each model answers one narrow question
4. Generative and agentic AI	Reads documents, summarizes histories, drafts letters, plans multi-step work	GenAI document processing, agents within governed workflows	Reliability, cost and oversight at scale

The shift from generation 3 to generation 4 changes who decides the path. In earlier generations, a designer drew every route a claim could take. An agent can choose its own next step toward a goal. In a regulated claims process, that freedom must be bounded: the workflow still owns the deadlines, approvals and audit trail, and the agent works inside them.

This matches a wider caution from Gartner, which predicts that more than 40% of agentic AI projects will be cancelled by the end of 2027 and notes that many use cases labelled agentic do not need agents at all (Gartner, 2025).

4. Pega as the Claims Automation Platform

Pega's main strength for financial claims is that AI runs inside a governed case. The case holds the deadlines, rules, approvals and audit trail, and AI models and agents do specific work within it. The table summarizes the capabilities most relevant to claims.

Pega capability	What it is	Role in financial claims
Case management and rules	Pega's core workflow engine: case types, stages, steps, SLAs and decision rules	The system of record for each claim: tracks regulatory clocks, enforces approvals, keeps the audit trail
Pega Blueprint	A generative AI design tool that drafts workflows, data models and personas from a business description and legacy documents (Pega Academy, n.d.-b)	Speeds up designing or modernizing a claims or dispute process; output is a starting draft that experts must review and approve
Process AI and adaptive models	Predictive models embedded in workflow steps	Scores complexity and outcome likelihood to route claims and pick straight-through candidates
Predictable AI agents (Infinity 26)	Agents designed with heavy AI reasoning at design time and efficient execution at runtime, orchestrated inside governed workflows	Handles judgment-based steps such as evidence review while the case controls what the agent may do
Model Context Protocol (MCP) support	Pega applications can act as MCP servers; Pega can also orchestrate third-party agents	Lets external AI assistants start or advance a claim or dispute through the governed Pega workflow instead of custom integrations
Smart Dispute Agentic Automation	Packaged application for card disputes, chargebacks and fraud claims (Pegasystems, n.d.-c)	Prebuilt Visa, Mastercard and Amex network flows; SLAs for Reg E, Reg Z, Canada's C-86 and UK Section 75
Smart Investigate Agentic Automation	Packaged application for payment exceptions and investigations	Automates research and resolution of failed or misrouted payments
Pega Claims for Insurance and Life	Packaged claims applications for insurers	End-to-end flow from first notice of loss to payment, with triage, fraud checks and correspondence

4.1 Agentic capabilities in Infinity 26

Pega announced Infinity 26 at PegaWorld in June 2026 and made it generally available on 14 July 2026 (Pegasystems, 2026b). Three changes matter most for claims teams:

- **Pricing tied to outcomes.** Pega said clients would pay a flat price per completed case instead of metered token charges (Pegasystems, 2026a). Industry coverage advises buyers to confirm how "completed case" is defined in the contract (Carter, 2026).
- **Open agent connectivity.** Through MCP support, agents built on other platforms can call Pega workflows, with dispute resolution and claims processing given as examples (Pegasystems, 2026c).
- **AI-assisted build.** The new Infinity Studio pairs a client's chosen coding agent with Blueprint's best practices (destinationCRM, 2026).

4.2 Banking disputes: Smart Dispute

Smart Dispute encodes network and regulatory rules directly in the workflow. Pega refreshes its Visa and Mastercard rule content twice a year, in April and October (Pegasystems, 2023). The application includes Reg E logic for provisional credit, customer attestation and case finalization (Pega Academy, n.d.-a). Its Process AI models predict whether a chargeback is likely to be accepted under network validation rules, so analysts can focus on disputes worth pursuing. Pega's product sheet reports client results of an 86% reduction in average resolution time and a 50% reduction in staff (Pegasystems, n.d.-a). These are vendor-reported figures and should be tested against an institution's own baseline.

4.3 Insurance claims

For insurers, Pega's claims optimization uses machine learning to flag claims that look simple at first notice but are likely to become complex, so they can go straight to an experienced adjuster. The same models predict claim outcomes to find candidates for straight-through processing, and teams can test models against live production data before relying on them (Pegasystems, n.d.-b). This fits the direction McKinsey set out in its *Claims 2030* research: for simple, predictable claims, the technology for full straight-through processing already exists (McKinsey & Company, n.d.-a).

5. The AI-Enabled Claims Lifecycle on Pega

Every financial claim, whether an auto loss or a disputed card charge, passes through the same six stages. The table shows the author's reference design: what AI does at each stage, which Pega capability carries it, and what stays with people.

Stage	What AI does	Pega capability	Human role
1. Intake (first notice of loss or dispute)	Captures the claim from any channel, reads uploaded documents and photos, pre-fills case data, checks eligibility	Case creation through web, mobile, chat or MCP-connected agents; GenAI document processing	Help customers who need it; confirm unclear facts
2. Triage and routing	Scores complexity, fraud risk and straight-through eligibility; starts regulatory clocks such as Reg E deadlines	Process AI models, decision rules, SLAs	Set thresholds; review high-risk flags
3. Investigation	Gathers policy, transaction and merchant data; summarizes claim and call history; drafts network chargeback filings	Agents within governed steps; integrations to card networks and core systems	Weigh conflicting evidence; contact parties when needed
4. Decision	Recommends approve, deny, partial payment or representment, with reasons	Decision rules plus AI recommendations, with approval steps by value and risk	Make or confirm any adverse or high-value decision

Stage	What AI does	Pega capability	Human role
5. Settlement and correspondence	Calculates payment or provisional credit, posts adjustments, drafts compliant letters	Financial adjustment steps, correspondence templates with regulatory language	Sign off on exceptions and sensitive letters
6. Recovery and learning	Pursues subrogation or chargeback recovery; feeds outcomes back to models	Recovery case types, adaptive learning, reporting	Review model performance and fairness

Two design principles run through the table. First, **the case is in charge, not the model**. AI proposes, extracts and drafts, but the Pega case decides what happens next under rules people can inspect. Second, **autonomy grows with evidence**. A low-value, low-risk card dispute with clear evidence can move straight through; a contested life claim cannot. Thresholds should be set per claim type and widened only as measured accuracy supports it.

6. Evidence: Adoption and Productivity

The evidence supports real gains from AI in claims-like work, but most organizations have not yet turned AI use into enterprise results. The gap is in workflow redesign, which is exactly the layer a platform such as Pega provides.

6.1 Adoption

McKinsey's 2025 global survey of 1,993 respondents found that 88% of organizations use AI in at least one business function. However, only 23% are scaling agentic AI anywhere, and no more than 10% are scaling agents in any single function. A small group of high performers, about 6% of respondents, stood out mainly because they had redesigned workflows rather than simply adding tools (McKinsey & Company, 2025; IDM, 2025).

For claims specifically, McKinsey estimates that more than half of current claims activities could be automated by 2030, which would remove some roles, create new digital ones and change the rest (McKinsey & Company, n.d.-b). Gartner expects at least 15% of routine work decisions to be made autonomously by agentic AI by 2028 (Gartner, 2025).

6.2 Productivity research relevant to claims

Study	Setting	Finding	Relevance to claims
Brynjolfsson, Li & Raymond (2025)	5,179 customer support agents with an AI assistant	Issues resolved per hour up 14% on average and 34% for novice and lower-skilled agents	Mirrors claims contact centers and dispute intake teams
Dell'Acqua et al. (2023)	758 BCG consultants, randomized	Inside AI's capability: 12.2% more tasks, about 25% faster, higher quality. Outside it: 19 percentage points less likely to be correct	Warns against applying AI to claim decisions it cannot handle reliably

Study	Setting	Finding	Relevance to claims
Noy & Zhang (2023)	453 professionals on writing tasks	Time down 40%, quality up 18%	Applies to claim letters, denial explanations and case summaries

The research yields one lesson that matters most for claims: AI capability is uneven. Dell'Acqua and colleagues describe a "jagged technological frontier," in which tasks of similar apparent difficulty fall on opposite sides of what AI does well. In claims, reading a repair estimate may be well within that frontier, while judging a contested liability may not be. Each step must be tested, not assumed.

7. Risk, Compliance and Governance

In financial claims, an AI error is not just a wrong answer: it can be an unfair denial, a missed statutory deadline or a payment to a fraudster. Governance must therefore control what automated steps are allowed to do, and prove afterwards what they did.

7.1 Regulatory landscape

Framework	Who it applies to	What it requires for AI in claims
NAIC Model Bulletin on the Use of AI Systems by Insurers	US insurers in states that adopt it	A written AI program that manages risks of inaccurate or unfairly discriminatory decisions, including oversight of third-party AI vendors
Regulation E (12 CFR 1005.11)	US banks and credit unions handling electronic fund transfer errors	Investigation, provisional credit and notice deadlines, whether work is done by people, software or vendors
EU AI Act	Providers and deployers of AI in the EU	Transparency duties from 2 August 2026; high-risk obligations from 2 December 2027 (Annex III) and 2 August 2028 (Annex I)
NIST AI Risk Management Framework	Voluntary, any organization	Four functions to govern, map, measure and manage AI risk

The NAIC adopted its model bulletin on 4 December 2023 (National Association of Insurance Commissioners, 2023). By March 2025, 24 states had adopted it with little or no change (Quarles & Brady, 2025), and some states, including Colorado and New York, have taken their own approaches (JD Supra, 2024). The bulletin names the NIST framework as one standard an insurer's AI program may rely on (Optro, 2026).

The EU AI Act timeline reflects the Digital Omnibus on AI, which the Council of the EU approved on 29 June 2026 (Future of Privacy Forum, 2026; Secure Privacy, 2026). Annex III's high-risk list includes credit scoring and risk pricing in life and health insurance, so European insurers and lenders should check whether any claims-adjacent models fall within it. The NIST AI RMF organizes risk work into Govern, Map, Measure and Manage (National Institute of Standards and Technology, 2023).

7.2 AI-specific risks

- **Over-reliance.** Reviewers stop checking AI output. In the BCG experiment, consultants using AI on a task beyond its ability did worse than those without it (Dell'Acqua et al., 2023).
- **Prompt injection.** A claim document or email can carry hidden instructions to an agent that reads it. OWASP ranks prompt injection as the top risk for LLM applications (Coralogix, 2025).
- **Excessive agency.** An agent with more tools or permissions than it needs can turn a small error into a wrong payment (Cloudflare, n.d.).
- **Unfair outcomes.** Fraud and triage models can disadvantage groups of claimants if they are not tested for bias.
- **Vendor dependence.** Regulators hold the institution responsible even when a vendor's system misses a deadline, as the 2024 Reg E findings showed.

7.3 Controls to build into the Pega workflow

1. Keep regulatory deadlines in case SLAs and rules, not in AI prompts, so they are enforced the same way every time.
2. Give each agent only the case actions and data it needs; route payments, denials and external letters through approval steps.
3. Treat all claimant-supplied documents as untrusted input, and keep data access controls in the platform.
4. Record every AI recommendation, the evidence used and the human decision in the case audit trail.
5. Route low-confidence or high-impact cases to people, and test models for accuracy and bias before and after each change.
6. Maintain an inventory of AI uses by claim type, owner and risk tier, mapped to the NAIC bulletin and NIST functions.

8. Implementation Roadmap on Pega

The fastest safe path is to start with one high-volume, well-regulated claim type, prove results in shadow mode, and widen AI autonomy only as evidence builds. The five phases below are the author's framework, aligned with the NIST functions.

1. **Govern: set ownership and guardrails.** Name an accountable executive for claims AI. Write the AI program the NAIC bulletin expects, or its banking equivalent, and decide which decisions will always need a person, such as claim denials or fraud referrals. *Exit test:* an approved policy, a risk tier for each claim type and a named owner for each automated workflow.
2. **Map: pick the first claim type and measure the baseline.** Good first candidates are low-value card disputes, simple auto glass or property claims, or payment investigations. Record today's

cycle time, cost per claim, regulatory exceptions, leakage and customer satisfaction. *Exit test:* a baseline signed off by operations, finance and compliance.

3. **Design in Blueprint, build on the packaged app.** Use Pega Blueprint to draft the redesigned case lifecycle from current procedures, then refine it with adjusters, dispute analysts and compliance staff. Start from Smart Dispute, Smart Investigate or Pega Claims rather than a blank application, and decide step by step whether rules, Process AI or an agent fits best. *Exit test:* a reviewed design showing every human checkpoint, SLA and system the automation can touch.
4. **Measure: run in shadow mode, then supervised.** Let AI score and recommend while people still decide, and compare its calls with theirs. Move to supervised automation, where AI acts and people review samples, only for steps that match or beat the baseline. *Exit test:* agreed accuracy, deadline compliance and fairness thresholds met, with no unresolved high-severity issues.
5. **Manage: scale by claim type and keep monitoring.** Expand straight-through processing one segment at a time. Re-test models after every rule, network or product change, especially after the twice-yearly card network rule releases. Report business outcomes such as cycle time, cost per claim, regulatory exceptions and complaints, not just usage. *Exit test:* sustained results over at least two reporting cycles before the next expansion.

On cost, Infinity 26's per-completed-case pricing makes it easier to model unit economics, but the business case should still be built on the institution's own measured baseline rather than vendor benchmarks.

9. Workforce Implications

AI will change claims and dispute roles more than it removes them in the near term. People will handle fewer routine steps and more judgment, exceptions and difficult conversations.

- **From processing to supervising.** Adjusters and dispute analysts will spend less time keying data and searching systems, and more time reviewing AI recommendations, handling escalations and talking with claimants. Performance measures should reward decision quality, not only throughput.
- **New roles.** Claims operations will need people who tune decision rules and thresholds, monitor model performance and fairness, and own Pega case designs. McKinsey's claims talent research expects new digital roles to appear as some current roles are eliminated (McKinsey & Company, n.d.-b).
- **Faster ramp-up for new staff.** In the customer support study, AI helped novice agents most (Brynjolfsson et al., 2025). Claims organizations can use AI guidance in Pega to train new adjusters faster, but should also protect the learning that once came from routine cases.
- **Wider labor-market context.** The World Economic Forum projects that between 2025 and 2030, 170 million jobs will be created and 92 million displaced worldwide, a net gain of 78 million, and

lists upskilling as employers' top workforce strategy (World Economic Forum, 2025). These figures cover all drivers of change, not AI alone.

10. Conclusion

Financial claims are where automation pays off most and where mistakes cost the most. AI can now read the documents, summarize the history, score the risk and draft the response. A workflow platform such as Pega supplies what AI lacks on its own: enforced deadlines, network and regulatory rules, approval steps and a complete audit trail.

The winning pattern is to keep the case in charge and let AI work inside it. Institutions that start with one claim type, prove results in shadow mode, govern under the NAIC bulletin, Regulation E and the NIST framework, and widen autonomy only as evidence builds will settle honest claims faster, catch more fraud, and be able to show regulators exactly how every decision was made.

References

Benton, K. J. (2025, October 28). Error resolution procedures under the Electronic Fund Transfer Act and Regulation E. *Consumer Compliance Outlook*, Third Issue 2025. Federal Reserve System. https://www.consumercomplianceoutlook.org/articles/2025/third-issue-2025/error-resolution-procedures-under-the-electronic-fund-transfer-act-and-regulation-e

Brynjolfsson, E., Li, D., & Raymond, L. R. (2025). Generative AI at work. *The Quarterly Journal of Economics, 140*(2), 889–942. https://doi.org/10.3386/w31161

Carter, R. (2026, September 14). Pegasystems has put a price on agentic customer service. Buyers need the fine print. *CX Today.* https://www.cxtoday.com/?p=992594

Center for Risk and Responsibility. (2024). *Misinformation about insurance fraud.* Rutgers Law School. https://crr.rutgers.edu/wp-content/uploads/sites/22/Misinformation-About-Insurance-Fraud-11-2024.pdf

Cloudflare. (n.d.). *What are the OWASP Top 10 risks for LLMs?* https://workers.cloudflare.com/learning/ai/owasp-top-10-risks-for-llms

Coalition Against Insurance Fraud. (2022). *The impact of insurance fraud on the U.S. economy.* https://insurancefraud.org/wp-content/uploads/The-Impact-of-Insurance-Fraud-on-the-U.S.-Economy-Report-2022-8.26.2022.pdf

Coralogix. (2025). *OWASP Top 10 for LLM applications: What you need to know.* https://coralogix.com/ai-blog/owasp-top-10-for-llm-applications/

Dell'Acqua, F., McFowland, E., III, Mollick, E., Lifshitz-Assaf, H., Kellogg, K. C., Rajendran, S., Krayer, L., Candelon, F., & Lakhani, K. R. (2023). *Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality* (Harvard Business School Working Paper No. 24-013). https://ssrn.com/abstract=4573321

destinationCRM. (2026, July 15). *Pega releases Infinity 26.* https://www.destinationcrm.com/Articles/CRM-News/CRM-Across-the-Wire/Pega-Releases-Infinity-26-175738.aspx

Electronic Code of Federal Regulations. (n.d.). *12 CFR 1005.11 – Procedures for resolving errors.* https://www.ecfr.gov/current/title-12/chapter-X/part-1005/subpart-A/section-1005.11

Federal Reserve System. (2026). Top Federal Reserve System violations in 2024: Regulation E error resolution requirements. *Consumer Compliance Outlook*, First Issue 2026. https://www.consumercomplianceoutlook.org/articles/2026/first-issue-2026/top-federal-reserve-system-violations-in-2024-regulation-e-error-resolution-requirements

Future of Privacy Forum. (2026). *The AI Act implementation timeline: What changes under the AI Omnibus?* https://fpf.org/blog/the-ai-act-implementation-timeline-what-changes-under-the-ai-omnibus

Gartner. (2025, June 25). *Gartner predicts over 40% of agentic AI projects will be canceled by end of 2027* [Press release]. https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027

IDM. (2025). *McKinsey: Few organisations fully scaling AI despite uptake.* https://idm.net.au/node/15415

JD Supra. (2024). *Rhode Island adopts NAIC AI model bulletin.* https://production.jdsupra.com/legalnews/rhode-island-adopts-naic-ai-model-1349751

McKinsey & Company. (2025, November). *The state of AI in 2025: Agents, innovation, and transformation.* https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai

McKinsey & Company. (n.d.-a). *Claims 2030: Dream or reality?* https://www.mckinsey.com/capabilities/operations/our-insights/claims-2030-dream-or-reality

McKinsey & Company. (n.d.-b). *Claims 2030: A talent strategy for the future of insurance claims.* https://www.mckinsey.com/industries/financial-services/our-insights/claims-2030-a-talent-strategy-for-the-future-of-insurance-claims

National Association of Insurance Commissioners. (2023). *NAIC model bulletin: Use of artificial intelligence systems by insurers* [Adoption map]. https://content.naic.org/sites/default/files/legal-adoption-map-ai-model-bulletin.pdf

National Institute of Standards and Technology. (2023). *Artificial intelligence risk management framework (AI RMF 1.0)* (NIST AI 100-1). https://doi.org/10.6028/NIST.AI.100-1

Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science, 381*(6654), 187–192. https://doi.org/10.1126/science.adh2586

Optro. (2026). *NAIC AI model bulletin.* https://optro.ai/blog/naic-ai-model-bulletin

Pega Academy. (n.d.-a). *Regulation E in Smart Dispute Agentic Automation.* https://academy.pega.com/node/91651

Pega Academy. (n.d.-b). *Pega GenAI Blueprint for efficient application design and development.* https://academy.pega.com/node/89491

Pegasystems. (2023, September 26). *Pega infuses AI into Pega Smart Dispute to streamline chargeback processes* [Press release]. https://pega.com/node/427541

Pegasystems. (2026a, June 8). *Pega eliminates 'AI token tax' with more efficient way to build and run agentic workflows* [Press release]. https://www.pega.com/node/467551

Pegasystems. (2026b, July 14). *Pega Infinity 26 now available to deliver predictable outcomes with predictable AI costs* [Press release]. https://www.businesswire.com/news/home/20260714981812/en/

Pegasystems. (2026c, June 8). *Pega powers AI agents to reliably drive mission-critical work* [Press release]. https://www.businesswire.com/news/home/20260608601073/en/Pega-Powers-AI-Agents-to-Reliably-Drive-Mission-Critical-Work

Pegasystems. (n.d.-a). *Pega Smart Dispute Agentic Automation* [Product sheet]. https://pega.com/system/files/resources/pdf/pega-pega-smart-dispute-agentic-automation.pdf

Pegasystems. (n.d.-b). *Pega claims optimization at a glance* [Video]. https://www.pega.com/node/151121

Pegasystems. (n.d.-c). *Pega Smart Dispute Agentic Automation for issuers.* https://pega.com/industries/financial-services/smart-dispute

Quarles & Brady. (2025). *Nearly half of states have now adopted NAIC model bulletin on insurers' use of AI.* https://www.quarles.com/newsroom/publications/nearly-half-of-states-have-now-adopted-naic-model-bulletin-on-insurers-use-of-ai

Secure Privacy. (2026). *EU AI Act Digital Omnibus: The new high-risk AI deadlines after Council approval.* https://secureprivacy.ai/blog/eu-ai-act-digital-omnibus-the-new-high-risk-ai-deadlines-after-council-approval

World Economic Forum. (2025, January 8). *Future of Jobs Report 2025: 78 million new job opportunities by 2030 but urgent upskilling needed to prepare workforces* [Press release]. https://www.weforum.org/press/2025/01/future-of-jobs-report-2025-78-million-new-job-opportunities-by-2030-but-urgent-upskilling-needed-to-prepare-workforces/

About this paper

Author: Varun Vemulapalli. **Published:** November 2026.