Access Partnership

AETHRYX

November, 2026

The Cost of an Untraceable Decision

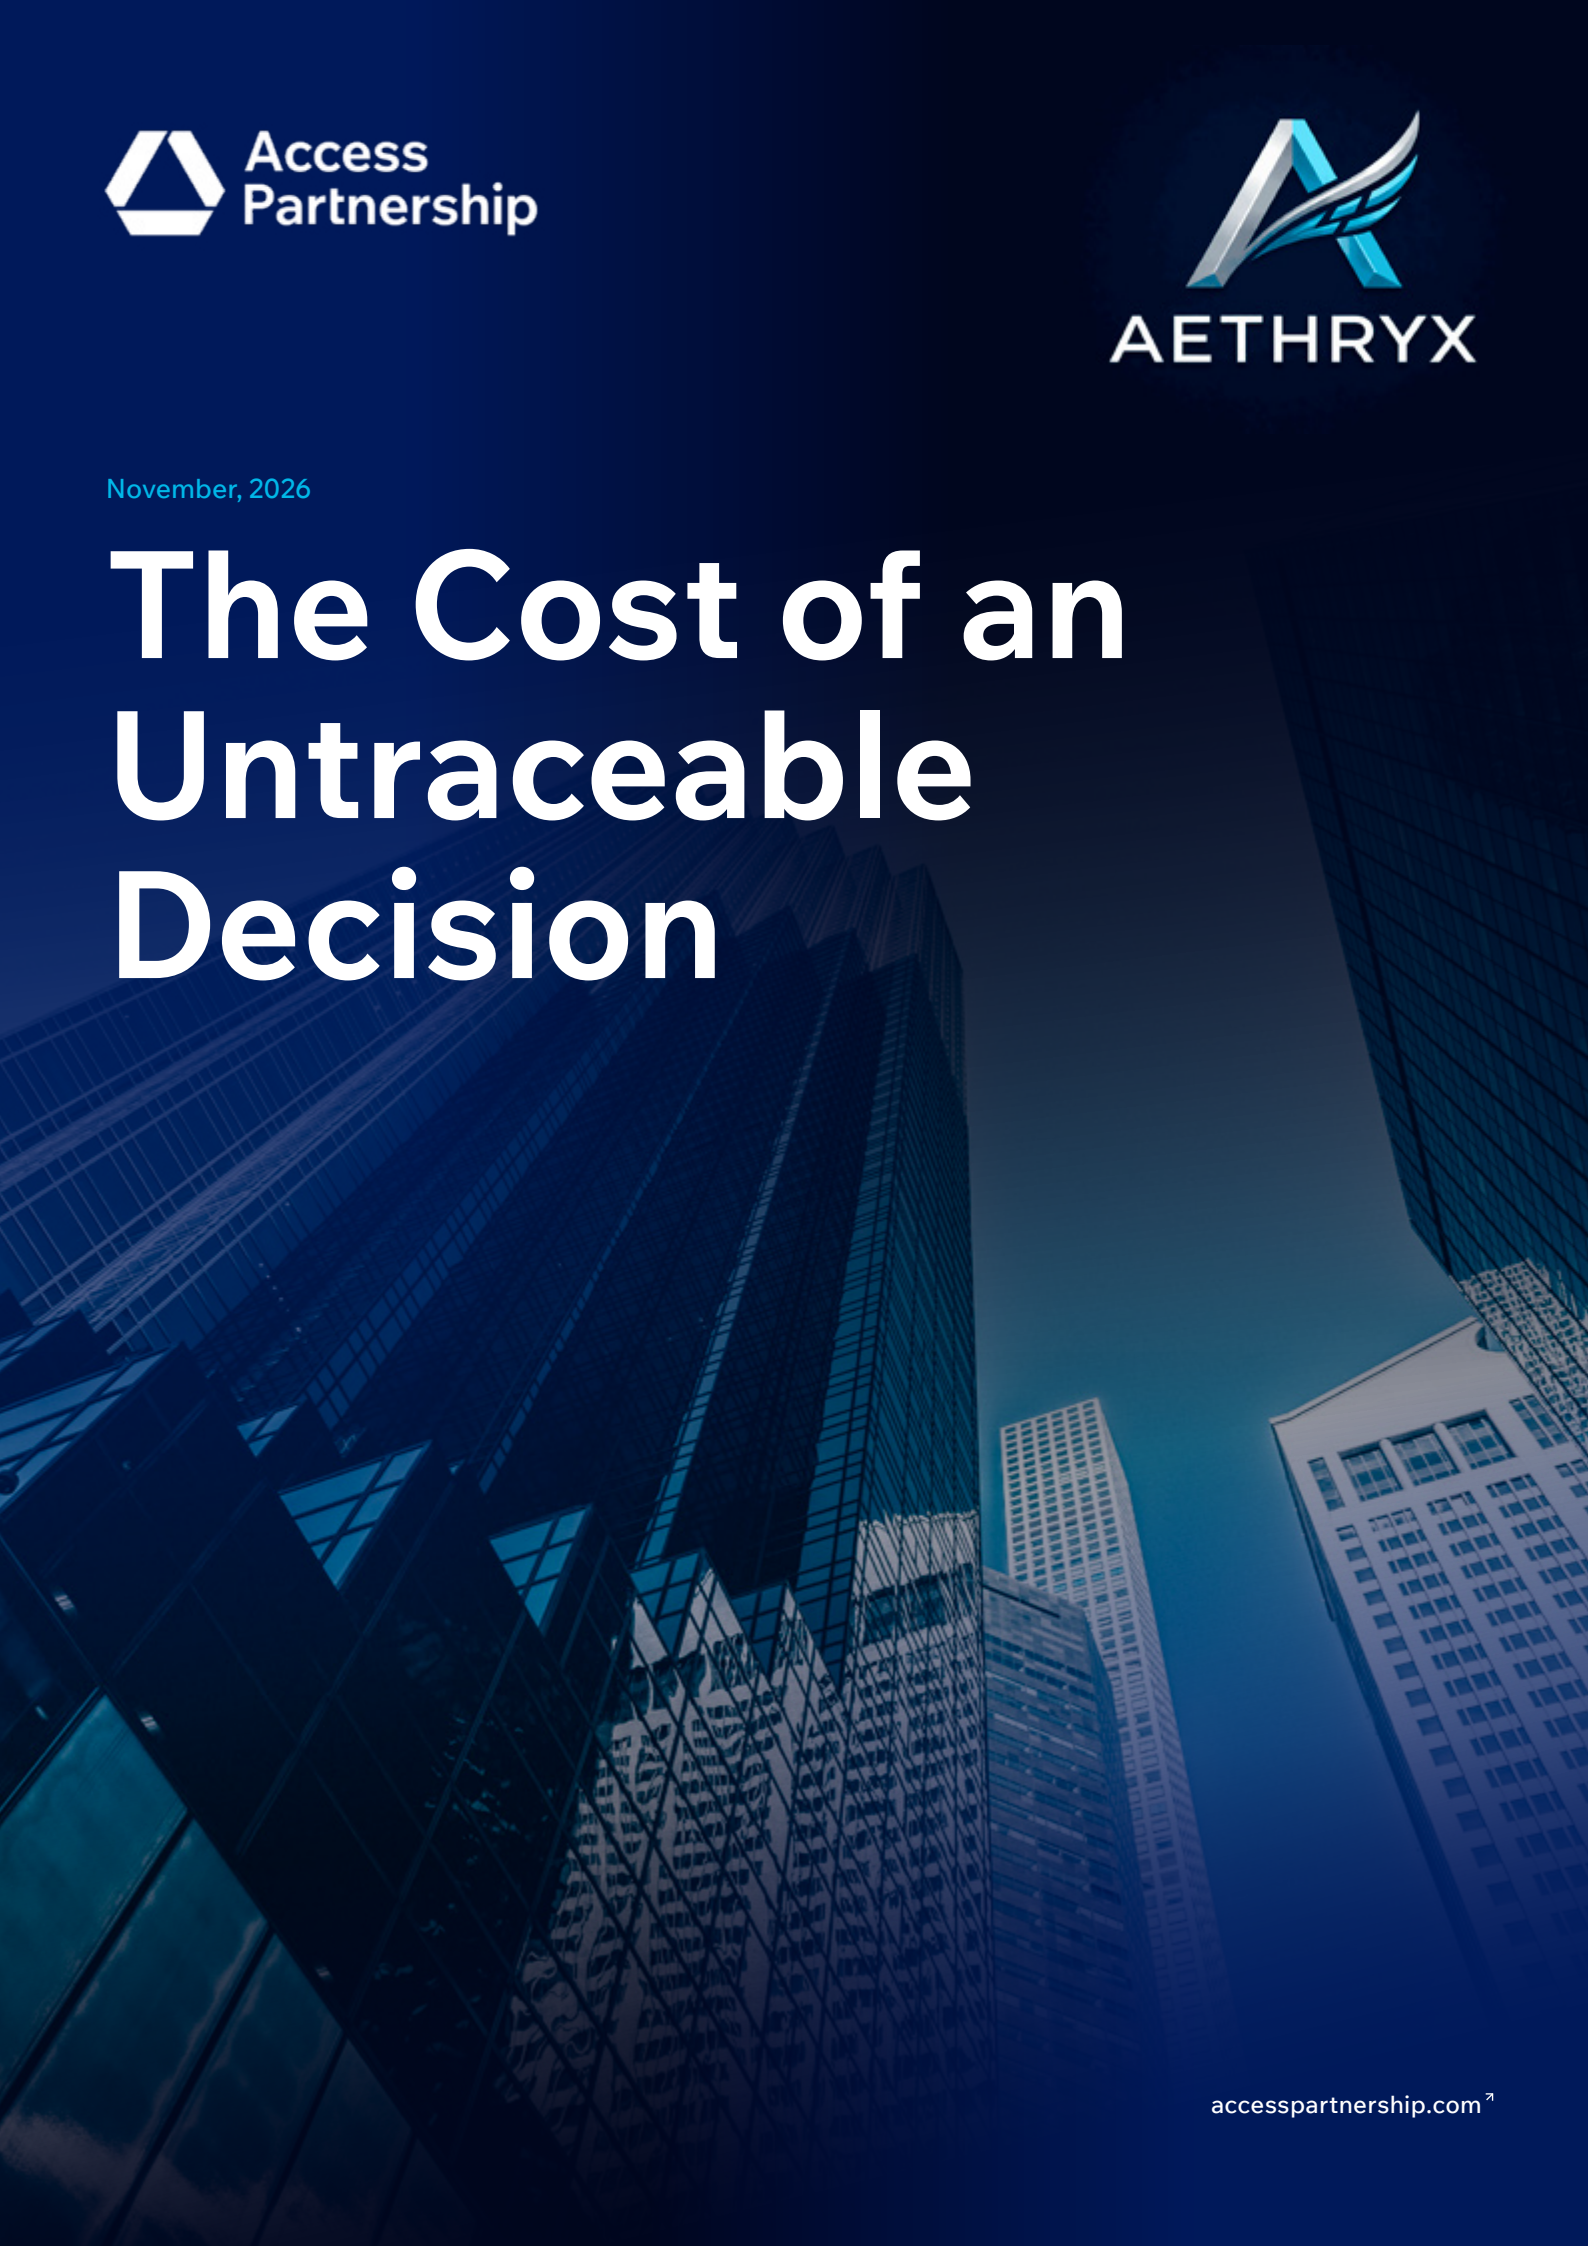

accesspartnership.com

Authors



Nada Aboserie

Deputy Director,
AI Advisory & Data Protection Officer,
Access Partnership | AI Governance
& Technology Policy Leader

Nada Aboserie is Deputy Director of AI Advisory and Data Protection Officer at Access Partnership, a global frontier consultancy, where she leads an international team and has designed the controls that govern the firm's intetnal use of AI.

With over a decade of experience working between technology companies and governments, she began her career at Microsoft and has since advised Google, Microsoft, AWS, Apple and Meta as well as governments in the Middle East and Asia Pacific regions. She built the government affairs strategy behind major hyperscale cloud build-outs and has trained more than 100 regulators on AI.

She is AIGP and CIPP/E certified and a co-author of Signals in the Noise, research on detecting AI loss of control.

An Aspen UK Rising Leaders Fellow and member of techUK's AI Forum, she has lived in nine countries and speaks five languages. She is driven to put people at the centre of the AI revolution, with a real sense of control over the technology shaping their lives.



Dr. Brindha Jeyaraman

Founder & CEO,
Aethryx | AI Governance & Agentic AI Leader

Dr. Brindha Jeyaraman is the Founder and CEO of Aethryx, a Singapore-based AI technology company building runtime governance solutions for enterprise AI and autonomous agents.

With over 19 years of experience in Artificial Intelligence, Machine Learning, and enterprise technology, she has held senior leadership roles at Google Cloud, UOB, the Monetary Authority of Singapore (MAS), and A*STAR.

She holds a Doctor of Engineering in Artificial Intelligence from Singapore Management University (SMU) and is an accomplished author, researcher, and speaker. Her work focuses on responsible AI, agentic AI, temporal knowledge graphs, and AI governance in regulated industries.

Recognised among the Top 50 Women Tech Leaders in Asia (2025), she is passionate about bridging AI innovation, enterprise adoption, and responsible governance to enable organisations to deploy AI safely, transparently, and at scale.

Contents

The 45-day Problem 4

01. When AI Takes Consequential Decisions 5

Regulators Are Converging on the Same Expectation 6

How Fast Automated Failure Can Move 6

02. Forty-five Days: A Hypothetical Credit Agent Scenario 7

The Lender 7

The Fault 7

How 45 Days Pass By 7

03. What 45 Days Cost 8

Inputs 8

The Cost 9

What Drives Each Part of The Cost 10

04. Runtime Control Architecture 11

05. The Same 45-days, Governed Differently 12

Isolating the Value of Evidence 12

Prevention is a Separate Case 13

06. A Practical Path to Traceable Decisions 14

07. Five Questions for The Next AI Review 15

References 16

This paper is written for board members and executives whose firms are deploying AI agents and who need to understand what a control failure would cost them. It is also for the people in risk, compliance, legal and technology who would be asked, after such a failure, to explain what the agent did and why nobody caught it sooner.

The 45-Day Problem

For **45 days**, an AI credit agent kept approving loans it should have declined. Nothing crashed and no alarm went off. The model had been tested and approved, and the agent stayed within the limits it had been given the whole time. What had changed was one of the data feeds it relied on, which had been restructured without notice. By the time the lender found the problem, the agent had taken **40,500 decisions**, and around **3,645** of the **loans** it approved should have been refused.

On the assumptions in this paper, the incident cost **US$19.8 million**. That is roughly **US$259,000** for every day the problem went unnoticed.

The lender is hypothetical and the figure is illustrative, but the scenario is a plausible one, and we have worked it through in detail because the value of a control is much easier to judge once you can put a price on its absence. The costing covers credit losses, customer remediation, investigation and business interruption. We have left out regulatory fines, litigation and reputational damage, since any figure for those would depend on how a particular regulator or market reacted, and we could not defend the assumptions.

Most AI governance concentrates on whether a system is safe and compliant at the moment it is approved. An agent does more than produce an answer, though. It gathers data, calls tools, applies rules and acts on the result, and the systems it depends on go on changing after sign-off even when nobody touches the agent itself. The questions a leadership team needs to answer are therefore whether it can see what its agent is doing now and stop it when something goes wrong, and whether it could later prove what happened.



01

When AI Takes Consequential Decisions

Most financial institutions built their AI governance around scoring models. A model is validated before it goes live and reviewed periodically, and its output passes to a person or a rules engine that owns the decision. That arrangement puts the main control at the point of approval and assumes the system will behave in production much as it did in testing.



A credit agent breaks that assumption. It retrieves a bureau file, calls an income verification service, runs an affordability calculation, applies credit policy, sets a limit and price, and issues the decision, often with no person involved in any individual case. Each of those steps depends on something that can change after approval. Data vendors revise their feeds, internal tools get updated, prompts and configurations are edited, and model providers release new versions. The agent can remain exactly as it was validated while the environment around it shifts, so the questions governance has to answer stop being settled once and need answering continuously.

Question	Answered once, at approval	Needs answering at runtime
Authority: what can the agent decide or execute?	Delegated limits in policy and the deployment sign-off	Whether the agent should still hold that authority, given the quality of the inputs it is using today
Access: what data, systems and tools can it use?	Referral rules and a quality assurance sampling plan	Whether those integrations still return what they returned when they were approved
Approval: when must a human review or intervene?	Referral rules and a sampling plan for quality assurance	Whether conditions have changed enough that more decisions should go to a person
Audit: can consequential decisions be reconstructed?	A logging standard and a retention period	Whether the firm can show, for any single decision, which data, tools and policy version produced it



Regulators Are Converging on the Same Expectation

In the European Union, AI systems used to assess creditworthiness are classified as high-risk under Annex III of the AI Act. The Digital Omnibus on AI, in force since 27 July 2026, moved the date on which Annex III obligations apply from 2 August 2026 to 2 December 2027. The classification itself is unchanged, so lenders deploying credit agents in the EU will still have to meet the Act's requirements on risk management, record-keeping, logging and human oversight.

The Monetary Authority of Singapore consulted in November 2025 on Guidelines on AI Risk Management for all financial institutions, closing on 31 January 2026. The proposals place accountability with the board and senior management, expect firms to keep an inventory of their AI use and assess its materiality, and cover AI agents explicitly. Credit decisioning is among the higher-risk uses that would attract stronger controls, including closer human oversight.

The UK Financial Conduct Authority has said it will not introduce AI-specific rules and will rely on the Consumer Duty and the Senior Managers and Certification Regime instead. In practice a named senior manager remains accountable for credit outcomes produced by an agent, and the firm has to be able to show that those outcomes are good for customers rather than assume it.

What these regimes have in common is that each asks the firm to demonstrate, after the event, what its system did and who was responsible for it. That is as much a question of evidence as of control, and it is the question the scenario in Section 2 is built to test.

How Fast Automated Failure Can Move

The best-documented example of how quickly an automated system can generate losses comes from outside credit. On 1 August 2012, an order router at Knight Capital, trying to fill 212 customer orders, sent more than 4 million orders into the market in about 45 minutes. The firm traded over 397 million shares and lost more than US$460 million. The SEC found that Knight lacked adequate controls to stop erroneous orders reaching the market, and the case was its first enforcement action under the 2010 Market Access Rule.

Knight was a trading algorithm rather than an AI agent, and it plays no part in our numbers. We mention it for two reasons. The gap between machine speed and human detection was measured in minutes, and the firm had reviewed its controls against the rule earlier that year and judged them adequate, yet nothing acted at the point where the orders were actually generated. A credit agent fails more slowly, because losses only surface as customers start missing repayments, and that slowness is exactly what allows a fault to run for weeks.

02

Forty-five Days: A Hypothetical Credit Agent Scenario

The Lender

The Lender is a mid-sized consumer lender offering unsecured personal loans through digital channels, with an average balance of US$12,000. It receives 1,500 applications a day, every day. Nine months before the incident it deployed an AI credit agent to handle applications end to end within delegated authority. The agent decides 60% of applications on its own, and the remaining 40% fall outside its authority and go to human underwriters.

The deployment went through independent model validation, legal and compliance sign-off and approval by the credit risk committee. Ongoing oversight consists of weekly management information on volumes and approval rates, a monthly quality review of a 1% sample of agent decisions and a quarterly model performance review. For each decision the Lender records the outcome, score, limit, price and reason codes, but it does not keep the intermediate tool calls or the raw values the agent retrieved. A named senior executive owns consumer credit decisioning, including the agent. By most current standards, this is a reasonably governed deployment.

The Fault

On Day 0 the Lender's credit bureau releases a scheduled update to the structure of its data feed. The field the agent's affordability tool reads for existing monthly credit repayments is reorganised, and for applicants with more than one revolving facility the tool now picks up only the first. It raises no error. The value it returns is plausible, just too low.

The agent does what it was built to do with what it is given. Understated repayments make applicants look more affordable than they are, and some whose true debt would have failed the affordability test are approved. Those wrongly approved loans are the subject of this paper. In practice the same fault would also produce some limits above what policy allows, and some decisions resting on incorrect data without any change in outcome. We have left those effects out to keep the case simple, which means the figures in Section 3 understate the full cost.

How 45 Days Pass By

Day	What happens	What the Lender sees
0	Bureau feed update goes live overnight	Nothing. No errors are logged and the agent's response times are normal
6	Weekly report shows the agent's approval rate up about four points	The rise sits inside tolerance and coincides with a marketing campaign, which gets the credit
14	Monthly quality review of a 1% sample	Sampled decision match recorded score and reason codes, so pass. Reviewers cannot see the raw bureau values because they were never kept
30	Month-end credit risk committee	Approvals are up, which the committee reads as the campaign working
36	First repayments fall due on week-one loans	Early missed payments in the new cohort run higher than in earlier ones, though the sample is still small
41	A collections analyst escalates the pattern	Credit risk opens an investigation
45	The data team compares bureau and open banking figures for the affected applicants	An incident is declared and the Lender begins assessing its obligations to customers and its regulator

The incident then enters a second phase. Nothing in the Lender's records shows which of the 40,500 agent decisions were affected, so all of them have to be reviewed. Human underwriters carry the agent's workload for a further 30 days while the fault is fixed, the agent is re-tested and the risk committee agrees to restore its authority.

03

What 45 Days Cost

The incident costs the Lender about US$19.8 million. Some US$11.6 million comes from loans that should never have been made, and that part grows by roughly US$259,000 for each day the fault goes undetected. The other US$8.1 million comes from reconstructing what the agent did and running the business without it, and it depends far more on the quality of the Lender's records than on how quickly the fault was noticed.

Inputs

Input	Value	Basis
Share of agent decisions wrongly approved	**9%**	The fault affects applicants with more than one revolving facility (30% of agent decisions), and 30% of those would have failed affordability on their true debt
Lifetime default rate	**28%**	Applicants who should have failed affordability, well above an in-policy book
Loss given default	**75%**	Typical for unsecured consumer lending
Remediation per customer	**US$650**	Affordability reassessment, interest and fee refunds, forbearance
Customer contact per wrongly approved customer	**US$25**	Letters, calls and complaint handling
Lookback review	45 minutes per decision at **US$65** an hour, fully loaded	Internal file review
External forensic team	6 people for 50 days at **US$2,500** a day	Reconstructing decisions without retained tool outputs
Internal incident team	2,400 hours at **US$120** an hour	Credit risk, data, technology and compliance staff
Legal advice	**US$600,000**	Lump sum
Independent assurance review	**US$1,500,000**	Assumes the regulator or board requires one
Manual fallback	30 days at **US$30** per application	Time to fix, re-test and restore authority
Applicants lost during fallback	10% of applications; 55% approval rate and **US$900** lifetime margin per loan	Abandonment from slower turnaround; illustrative

Over 45 days the Lender receives 67,500 applications, of which the agent decides 40,500 on its own. At 9%, that gives 3,645 wrongly approved loans, about 81 a day, with total exposure of US$43.7 million.

The Cost

Cost line	Calculation	US$
Decision Losses		
Expected credit losses	43,740,000 × 28% × 75%	**9,185,400**
Customer remediation	3,6w 45 × 650	**2,369,250**
Customer contact	3,645 × 25	**91,125**
Subtotal		**11,645,775**
Investigation and Evidence		
Lookback review of every agent decision	40,500 × 0.75 hours × 65	**1,974,375**
External forensic reconstruction	6 × 50 × 2,500	**750,000**
Internal incident team	2,400 × 120	**288,000**
Legal advice	Lump sum	**600,000**
Independent assurance review	Lump sum	**1,500,000**
Subtotal		**5,112,375**
Business Interruption		
Manual underwriting during fallback	900 a day × 30 days × 30	**810,000**
Lost lending margin	1,500 × 30 × 10% × 55% = 2,475 loans × 900	**2,227,500**
Subtotal		**3,037,500**
Total		**19,795,650**

Credit losses are counted in full rather than as the difference from an in-policy loan, because these loans should not have been made at all.

What Drives Each Part of The Cost

The three groups of cost respond to quite different things, and that matters for where a firm puts its money.

Decision losses grow with time to detection. Each day the fault ran, the agent approved about 81 loans it should have declined, at roughly US$3,195 each. Holding everything else constant, catching the fault on Day 7 would have limited this group to US$1,811,565, on Day 14 to US$3,623,130, and on Day 30 to US$7,763,850, against US$11,645,775 at Day 45. Some remediation costs are fixed and defaults would not fall evenly across cohorts, so this is a simplification, but it gives the right order of magnitude.

Investigation costs depend on what the records can show. The lookback covers all 40,500 agent decisions rather than the 3,645 that went wrong, because nothing in the Lender's records tells them apart, and the external forensic work is needed for the same reason. Earlier detection would only reduce these costs in proportion to the number of decisions taken. Had the fault been found on Day 20, the Lender would still have had to review every agent decision since Day 0 and rebuild each one without the underlying data. The problem is the absence of evidence, and time has little to do with it.

Business interruption depends on how long the agent stays switched off, which turns on how quickly the Lender can satisfy itself and its risk committee that the fault is fixed, and on whether it has any option short of withdrawing the agent's authority entirely.

Section 5 reruns this case with the same lender, volumes, unit costs and Day 0 fault, and sorts each cost line by whether better controls would prevent it, reduce it, or leave it unchanged. A board weighing investment in runtime controls should know which savings depend on stopping bad decisions and which depend on being able to explain them.



04

Runtime Control Architecture

The core of the architecture is an enforcement point placed in the path to loan approval and disbursement. Monitoring tools that observe the agent can collect evidence, but they cannot reliably prevent an action unless that action is required to pass a control check, and the agent must not hold any alternative credential or route that lets it go around the check.

Before any commitment, the gate evaluates the agent's current authority, its permitted data and tool access, the approvals the case requires and whether the evidence is complete. It returns one of three outcomes: allow, refer or block. Risk can be reassessed whenever a dependency changes or a discrepancy appears. If a required check or evidence write fails, the defined safe outcome is referral or suspension of the affected autonomous action, with a named owner and a response deadline.

For this Lender, the controls would validate feed versions and the meaning of repayment fields, reconcile facility counts against repayment totals, and require adequate input quality before an affordability result is accepted. A schema check on its own would not catch this fault, because the value returned was well formed and merely wrong. The relevant test is whether the repayment figure includes every facility policy requires, and the most practical way to expose an omission like this one is to compare against facility totals parsed independently, or against another sufficiently independent source. Running the same faulty parser on both sides of the comparison would offer little protection, and a semantic change that no available check can observe remains a residual risk.

Where repayment information is missing, inconsistent or unverified, the agent's authority to commit should be suspended and the case sent, with its supporting evidence, to an authorised underwriter. The approval or rejection, any override rationale and the expiry of that approval are all recorded.

Traceability means linking each application to its sources and retrieval times, the permitted source snapshot, the parser and tool versions, the derived repayment figures, the affordability calculation, the policy version, the agent's authority at that moment and the final action. These records should be written as the workflow runs, protected against alteration, and access-controlled, so that when something goes wrong the firm can produce an indexed incident pack listing affected decisions, control outcomes, approvals and remediation status.

It is worth being clear about what such a record can and cannot establish. Retaining observable inputs, transformations, tool outputs, policy evaluations and actions, with provenance and version identifiers, allows the recorded execution to be reconstructed. It does not reveal a model's private reasoning, and it does not guarantee that a later run would produce the same output. Where retention of source data is restricted, the resulting evidential gap should be documented and permitted snapshots or retrievable references used in its place. Hashes prove integrity only for as long as the underlying evidence is still available. Throughout, personal data should be minimised, access restricted and retention periods set, and the firm should test bypass resistance, approval handling, missing evidence, changed source semantics and recovery before delegating live authority.

05

The Same 45-days, Governed Differently

The same Lender receives 1,500 applications a day, delegates 60% to the agent and meets the same bureau change on Day 0. What follows is a counterfactual to be tested, and no particular product is assumed to have caught this incident.

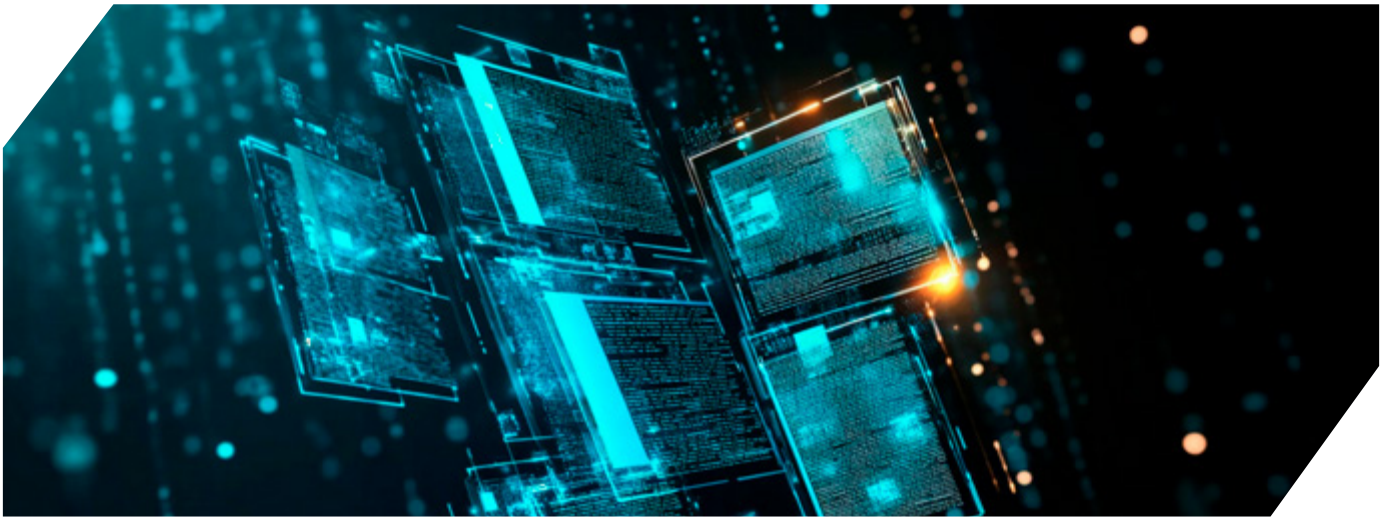

If a version or reconciliation check can observe the changed repayment representation, it catches it on Day 0 before any commitment, and the gate refers affected cases instead of approving them. If the check misses it, decisions continue and exposure accumulates as before, but the rest of the timeline changes. On Day 6 the rise in approval rates is reviewed alongside source versions, facility counts and repayment discrepancies, so the marketing campaign can explain the movement but cannot stand in for checking that inputs are complete. On Day 14 the quality review looks at retained source data and affordability calculations as well as scores and reason codes, and a single sampled error triggers a query for every other decision using the same source and parser. The Day 30 committee sees exception volumes, referral delays, unresolved data issues and authority changes alongside portfolio performance. When missed repayments surface between Days 36 and 41, they can be linked straight back to the original decisions, and by Day 45 the team can assemble the decision lineage, identify the affected cohort, validate the result and begin assessing customers. Unaffected activity resumes where dependencies are separated and the evidence supports it.

Isolating the Value of Evidence

To avoid crediting logging with prevention, the comparison below assumes that all 3,645 wrong approvals still happen. Decision losses, legal advice and independent assurance are unchanged. The only things that change are the effort to investigate and the length of manual fallback, on the assumption that complete, usable evidence and a validated repair allow a restart after five days instead of thirty.

The governed assumptions are ten minutes of review for each of the same 40,500 decisions, two external specialists for ten days at the original US$2,500 daily rate, 800 internal hours at US$120 and five days of fallback. They are sensitivity inputs and should not be read as measured productivity claims. Every decision stays in scope, so the result does not depend on narrowing the review.

Cost line (US$)	Baseline	Evidence only
Expected credit losses	9,185,400	9,185,400
Customer remediation	2,369,250	2,369,250
Customer contact	91,125	91,125
Review every agent decision	1,974,375	438,750
External forensic work	750,000	50,000
Internal incident team	288,000	96,000
Legal advice	600,000	600,000
Independent assurance	1,500,000	1,500,000
Manual fallback	810,000	135,000
Lost lending margin	2,227,500	371,250
Total	**19,795,650**	**14,836,775**
Illustrative reduction		**4,958,875**

The arithmetic behind the changed lines is:

Review 40,500 × 10/60 × US$65 = US$438,750

External Work 2 × 10 × US$2,500 = US$50,000

Internal Work 800 × US$120 = US$96,000

Fallback 900 × 5 × US$30 = US$135,000

Lost margin1,500 × 5 × 10% × 55% × US$900 = US$371,250

The reduction is **25.1%** of the baseline, before the cost of building and running the controls. Of it, **US$2,427,625** comes from lighter investigation and **US$2,531,250** from shorter interruption, and the **US$11,645,775** of decision losses is untouched. If fallback still lasted thirty days, the reduction would fall to **US$2,427,625**, since usable evidence on its own does not prove that a faster restart is possible.

Prevention is a Separate Case

At the baseline rate, 900 autonomous decisions a day produce 81 wrong approvals, each carrying US$2,520 of expected credit loss, US$650 of remediation and US$25 of contact cost, or US$258,795 a day in total. If an effective gate stopped further wrong approvals after two full days, and every referred case was then assessed correctly, 162 wrong approvals would remain and 3,483 would be avoided, a reduction in decision losses of 3,483 × US$3,195 = US$11,128,185.

That figure should not be added to the evidence savings above. Referring cases earlier changes workload, affected volumes, customer handling and interruption costs, so a prevention case has to recalculate all of those lines, including capacity for up to 270 multi-facility referrals a day, the resulting exception queues and any wrong approvals made by human underwriters. Catching the fault before the first commitment could avoid this fault's wrong approvals entirely, but the architecture has to demonstrate that in testing.

06

A Practical Path to Traceable Decisions

We propose a joint delivery model, for discussion, structured over six weeks. Timing and scope will depend on data access, integration readiness and how the lender owns its controls, so six weeks is a planning structure and should not be read as a guaranteed production deployment.



Before delivery starts, the lender's control owners and Access Partnership map obligations, materiality, policy and ownership, and agree authority boundaries, evidence retention, escalation routes and acceptance criteria. Access Partnership can lead this governance preparation. In the first two weeks the team maps the workflow, its source dependencies and action paths, instruments evidence capture and identifies where approval and disbursement must pass through enforcement gates. Weeks three and four translate the agreed policy into versioned checks, authority rules and approvals, integrate them with identity, workflow and incident processes, and define safe behaviour for referrals and outages. Week five replays permitted historical cases and tests feed changes, missing evidence, bypass attempts, approval expiry and semantic discrepancies, recording and fixing what fails. In week six the team demonstrates reconstruction of selected decisions and an incident response, reviews exceptions, operating cost and ownership, and seeks business and risk acceptance before any authority is expanded.

Acceptance should rest on operational outcomes. The lender should agree the evidence needed to reconstruct sampled decisions, test that out-of-policy actions cannot bypass the gate, measure how long it takes to assemble and validate evidence, confirm referral capacity and restart criteria, and test retention and access restrictions. Detection limits and residual risks should be recorded explicitly.

07

Five Questions for The Next AI Review

1. **Can we reconstruct the source inputs, transformations, policy checks and action behind a consequential decision made 60 days ago?**
2. **Would our monitoring detect a policy-relevant change in a dependency, or would it only tell us the infrastructure is healthy?**
3. **How quickly can we identify the potentially affected cohort, and how would we validate that we had found all of it?**
4. **Which controls can stop an out-of-policy action before it is committed, and is there any way for the agent to bypass them?**
5. **Who owns the decisions to restrict authority, approve exceptions, remediate customers and restore operation?**

The comparison in this paper is a conditional one. With the same wrong approvals, faster evidence review and a five-day fallback reduce an illustrative US$19.80 million incident to US$14.84 million, a saving of US$4.96 million before the cost of implementing and operating the controls. That is not a measured return on investment, and any claim about prevented credit losses requires controls that work in practice and a separate recalculation of referral and recovery costs.

References

Sources

- Aethryx, Executive AI Governance Series (Authority, Access, Approval and Audit framing)
- US SEC press release 2013-222, SEC Charges Knight Capital With Violations of Market Access Rule
- US SEC Administrative Order, Release No. 34-70694, In the Matter of Knight Capital Americas LLC
- Gibson Dunn, EU AI Act Omnibus Agreement: Postponed High-Risk Deadlines
- Usercentrics, EU AI Act Deal: Digital Omnibus Now in Force
- MAS media release, Guidelines for Artificial Intelligence Risk Management (13 November 2025)
- MAS Consultation Paper on Guidelines on AI Risk Management
- Bird & Bird, MAS Consults on Proposed Guidelines on AI Risk Management
- Freshfields, Navigating the New Regulatory Momentum: AI in UK Financial Services
- US Securities and Exchange Commission. SEC Charges Knight Capital With Violations of Market Access Rule. Press release 2013-222, 16 October 2013.
- US Securities and Exchange Commission. In the Matter of Knight Capital Americas LLC. Release 34-70694, 16 October 2013.

Access Partnership

accesspartnership.com

Need a problem solved?

Our dedicated experts, located around the world, are here to help.

Europe

London

Thomas House, 84 Eccleston Square
Office 5.06, 5th floor,
SW1V 1PX, London
United Kingdom

+44 20 3143 4900
london@accesspartnership.com

Brussels

8th Floor, Silversquare Europe
Square de Meeûs 35
B-1000 BrusselsBelgium

brussels@accesspartnership.com

North America

Washington DC

1300 Connecticut Avenue NW,
Suite 250
Washington, DC 20036
USA

+1 202 503 1570
washingtondc@accesspartnership.com

Asia

Bangkok

188 Spring Tower
11th Floor, Unit 106, Phayathai RoadThung
Phayathai, Ratchathewi, 10400
Thailand

+ 66 (2)-8216148

Hanoi

19th floor, Tower 1
Capital Place Building
No 29 Lieu Giai Street
Ngoc Khanh Ward, Ba Dinh District
Vietnam

Jakarta

Revenue Tower 21st Floor
Unit 104 SCBD Lot 13, Jl. Jend. Sudirman Kav.
52-53
Provinsi DKI Jakarta, 12190
Indonesia

+62 21 5020 0949

Manila

CG8ROCKWELL Level 21, 8 Rockwell,
Hidalgo Dr., Rockwell Center, Bgy.
Poblacion, 1210
Makati City, Metro Manila
Philippines

Kuala Lumpur

Common Ground Q Sentral
Level 39, Unit 39-02 (East Wing), 2A,
Jalan Stesen Sentral 2, Kuala Lumpur
Sentral, 50470
Malaysia

Singapore

12 Marina View
#11-20, Asia Square Tower 2
Singapore 018961

+65 8323 7855
singapore@accesspartnership.com

Follow us