

WHITEPAPER

Verified data in the AI era: The new trust layer for financial decisions



Executive Summary

Underwriting, audit, compliance and eligibility decisions are accelerating with the use of machines, and the cost of an unverified input is rising right along with them.

Artificial intelligence (AI) has done two things at once — it has made financial decisions faster, and it has quietly made the information behind those decisions easier to fake. As a result, the industry's trust problem has compounded. For years, the question was whether a document looked real, but the bigger question was really whether the information actually came directly from the party that owns the underlying record. Answering this requires a whole new kind of control — verification from the source.

The implication is straightforward; verification of authenticity must move upstream. It cannot sit at the end of the process — checking information after it has already entered the decision chain — because by the time it gets there, an unverified input may have already funded a loan, cleared a payment or become part of a system's downstream data.

Verification works better as a front-end control, establishing trust before information enters an automated decision, when the potential consequences of an error are still limited.

Organizations that build that gate now can treat verification as a strategic capability. Those that wait may find themselves retrofitting it later in response to fraud, regulatory expectations or reputational damage.

This paper examines why documents alone can no longer carry the weight the industry has placed on them, what automation can do to the cost of a bad decision and what a source-verified architecture looks like once in practice.

— *Shepley Smith, Chief Technology Officer,*
and *Jake Gutting, Chief Product Officer*



The unpriced risk in financial automation

AI is making financial decisions faster. It is also making the data behind those decisions easier to fake, creating a new tension between how fast decisions get made and trust in the inputs.^{1, 2, 3}

Five years ago, that tension mattered less, because two safeguards still stood between a bad document and a bad decision:

- Convincing forgeries were harder to produce⁴
- A person was usually somewhere in the chain reviewing the information

These safeguards have weakened at the same time, which is part of why the risk feels new even though document fraud itself is not.⁴

“AI needs good data” is true,⁵ but also incomplete.

Most AI systems can judge whether a document is:

- Structured correctly
- Internally consistent
- Visually consistent with an authentic document

None of these checks establish whether the information actually came from the party that

owns the underlying record. A document can look accurate, complete and authentic while its provenance remains unknown.^{2, 3, 6} A model trained to recognize patterns and formatting can assess what information looks like, but on its own, it cannot establish where information came from.

Once a decision is automated, that gap is no longer contained to a single transaction. An unverified input can move directly into a decision, into downstream systems that act on it and potentially into datasets or models that learn from it. This can carry a single bad input far past the point where anyone could have caught it just two short years ago.

Auditors, lenders, banks and other regulated organizations are already automating parts of these workflows and intending to move more and more humans out of the loop.⁷ What they need is a plain answer to a specific question: “What actually makes financial data trustworthy enough to support a decision no one is going to double-check?”



Where misplaced trust enters the system

In practice, unsupervised automation is making high-stakes calls no company would let an inexperienced employee make alone.

The gaps tend to show up in several familiar places:

- Borrower-provided income and asset documentation in lending ^{2,3}
- Self-reported financials in audit
- Benefactor-reported assets in government benefit eligibility determinations

The risk is not limited to outright fabrication. Understating, concealing or selectively presenting information that is otherwise legitimate creates the same underlying problem. That is because the decision still rests on something that was never independently checked.

This distinction mattered less when a human reviewer sat between the input and the outcome. It matters considerably more now as the input feeds directly into an automated decision, with fewer opportunities for a human to challenge or contextualize it.

From paper to digital: The limits of document-based verification

Document-based verification persists for a reasonable historical reason. For years, a scanned bank statement or uploaded PDF was treated as good enough. That's because forging a convincing one took real effort, and a human was somewhere in the review chain to catch what slipped through.⁴

Today's reality is different. Both conditions have disappeared, and not gradually. They disappeared almost simultaneously, which is partly why the exposure feels sudden even though document fraud itself has been around for decades.⁴

"Generative AI has closed the gap between a real document and a fabricated document almost down to the pixel level now," Shepley Smith,

*You
wouldn't
let an
intern set
your credit
risk model*

Jake Gutting,
Chief Product Officer,
Veriv



chief technology officer at Veriv, says, explaining that these convincing forgeries are being produced quickly and at low cost.

Each forged document will not necessarily evade every detection system; however, in this new reality, document inspection alone cannot establish provenance. A trained reviewer or optical character recognition system can identify inconsistencies, formatting anomalies and other signals. These checks, however, still cannot independently establish that the information came from the institution that owns the record.^{2, 6}

This is not a hypothetical concern for some future date.

“Even some of our least technically-savvy family members,” Gutting explains, “could go and generate a fake statement pretty quickly without expensive tools or training.”

Smith agrees the exposure is already growing, noting that ***“Any workflow that’s treating just the uploaded artifact as the truth is now bringing in a tremendous amount of risk.”***

Detection alone cannot win this race

The problem with relying on detection alone is structural — improving the very thing it is trying to catch.

“Every time an improvement in forgery detection gets folded into the next generation of document generation, that’s providing a signal back to the people who are forging it,” Smith explains.

Fraud can adapt as soon as detection improves, forcing institutions to keep updating the controls

used to distinguish genuine artifacts from fabricated ones.²

A system built to verify against the authoritative source approaches the problem differently. It asks whether the institution that owns the record actually supplied the information, rather than simply asking if the document appears real.²

The evidence: When document-level trust breaks down

Exposure is not merely theory.

In 2024, the Financial Crimes Enforcement Network,¹ part of the U.S. Department of the Treasury, formally warned that fraudsters were using generative AI to fabricate identity documents and circumvent identity verification and authentication methods at financial institutions. This Gen-AI trend is also tied to fraudulent accounts and check, credit card and loan fraud, with significant Bank Secrecy Act reporting implications for affected institutions.^{1, 9, 10}

Australia’s mortgage market offers a more concrete illustration of how document-level checks can fail.

In 2026, reports put suspected mortgage fraud involving AI-fabricated financial documents¹¹ at up to AU\$4 billion (about US\$2.8 billion), with the Commonwealth Bank of Australia alone investigating up to AU\$1 billion (about US\$695 million) in potentially fraudulent loans.¹²

How does this happen exactly? Some of the fraudsters involved reportedly ran real salary deposits through genuine accounts for months before applying, creating files in which the documents and transaction history appeared consistent from every angle a document reviewer would check. Unfortunately, document-level detection missed it.²



AI-fabricated financial documents up to

AU\$4 billion



The lesson is important. A document can pass a plausibility check while the underlying financial picture remains unverified. The question is therefore not simply whether the submitted artifact looks authentic, but instead whether the information behind the decision can be traced back to a trustworthy, verifiable source.²

Today, major consumer banks are warning consumers about AI-enabled fraud, which is a fair signal that the issue extends beyond isolated document-forgery cases and is now a mainstream concern.³

The economics of error: What changes once humans leave the loop

In the past, a poor manual decision might result in one bad loan or outcome, but the mistake stopped at the transaction because the workflow naturally limited how quickly and far that error could spread. Automation has changed that math.

The same unverified input can trigger automatic funding, approval or cleared payments without anyone reviewing it first.^{2, 4, 9, 13} The error can execute at scale before anyone even realizes it has

happened.

The reality is that automation does not always create new failure modes. Instead, it can amplify their impact.¹³ As Gutting explains, “It can be exponential in some of these scenarios, removing some of the checks and balances that have always existed.”

The result is a larger potential blast radius. An unverified input can affect the original decision, downstream systems, and — where data is reused for analysis or model development — potentially future decisions as well.

“The cost of an unverified input now scales with everything that’s built on top of it,” Smith says, explaining the shift in terms of scale.

The earlier verification occurs, the fewer opportunities for an unverified input to compound.

Looking authentic does not equal trustworthiness

Accurate-looking data can be internally consistent, properly formatted and entirely plausible.

Trustworthy data, however, is something much narrower. It is information whose provenance can be independently established and connected to the party that owns the underlying record.⁶

A well-structured fabricated statement and an authentic one can look identical to a system trained to recognize formatting and consistency rather than provenance, which is exactly why an AI model cannot make this distinction on its own.³ Pattern recognition can assess what the information looks like; it cannot, by itself, establish who supplied it.

“AI doesn’t just need data that looks right; it needs



data that's been confirmed by the party that actually owns the data and owns the record," Smith explains.

Trust, in other words, has to be established at the source rather than inferred from the artifact sitting on top of it.

"Trustworthy data," Smith states, "isn't just a quality that you inspect at the end of the process; it's something that you have to build for out of the gate."

3 tests: How to discern what makes data trustworthy

Three connected tests help determine whether information has a defensible basis for trust:

- 1  **Provenance**
Asks whether the data can be traced to its source, not merely to where it was last seen or handed off.
- 2  **Authentication**
Asks whether the party providing the data actually controlled the underlying record, or simply had a copy.
- 3  **Tamper Resistance**
Asks whether the record can be shown to have remained unaltered from where it left the authoritative source.

These three are connected controls — not isolated checkboxes. A weakness in any one of them can undermine confidence in the result.

"A perfectly authenticated response to the wrong question doesn't matter," Smith

explains. **"Authentication proves the information was properly controlled — it says nothing about whether it's the right information for the decision being made."**

That distinction makes authentication and authoritative sourcing complementary, not interchangeable.

"Authentication proves the party in the chain actually controls the record," Smith explains, further noting that "an authoritative source means you went to the party directly."

Verification: A gateway, not a checkbox

Verification has historically been sold and understood as a post-decision, quality-control step — something positioned after fraud already had room to occur.

This arrangement made more sense when a bad input caused one bad outcome. It makes considerably less sense once the cost of an unverified input increases every time it moves further through an automated workflow.

If that cost compounds, verification cannot remain a back-end, quality-control exercise. It has to become a pre-decision control, establishing trust before information even enters the workflow, while the potential blast radius is still small.

Historically, audit confirmation followed a similar arc. It was treated as a nice-to-have until fraud became



visible enough that regulation caught up and made it a requirement. There is little reason to expect lending, compliance and eligibility workflows to follow a different path.

Unverified data should therefore be treated as a risk exposure, not simply a data-integrity issue. Smith frames the stakes as a balance-sheet issue rather than a data-quality one, noting that “Unverified data isn’t just a data problem. It’s almost like an unrecognized liability sitting on a balance sheet that no one knows exists yet.”

Verification is a risk-reduction measure; not a guarantee against every possible failure. It cannot eliminate risk, but moving verification upstream can reduce the chance that unverified input compounds before a consequential decision is made.

Gutting explains the same gap from the responder side of the relationship, noting that “verification is still seen as a check-the-box task while it really should be strategic. It just hasn’t caught up yet.”

The architecture of a source-first system

Verification needs to function as infrastructure beneath a workflow, not another application layered on top of one.

The practical implication is that a workflow platform should be able to consume verified information without also being required to build and maintain the network of relationships and delivery methods required to obtain it in the first place.

“Infrastructure sits underneath every process and feeds it,” Smith explains, noting that “Ideally, nobody has to think about it.”

The source-first principle, then, is straightforward. A system should seek information directly from the authoritative source rather than treating intermediaries and uploaded artifacts as equivalent.²

“Every intermediary that sits between you and the authoritative source for that document is an unverified input, and that’s bringing in risk,” Smith notes.

Verification is still seen as a check-the-box task...

Jake Gutting,
Chief Product Officer,
Veriv



Not all authoritative sources are accessible through the same technology.

A verification layer therefore needs to work across:



It should also do so without requiring the organization on the other end to overhaul its existing technology stack simply to participate.

Veriv exists to own that layer. It is a network that connects organizations requesting financial information directly to the banks, credit unions, employers and government platforms that actually hold the record. Veriv then routes each request to whichever channel that source can respond through, rather than asking every institution to adopt the same one. Maintaining these connections — one bank at a time, one delivery method at a time — is Veriv's rather than the requestor's or the source's.

"We can do that at scale," Smith says of Veriv's abilities. ***"We can do that across multiple mediums, whether that's through an API or through a fax-based system or through users interacting with our portal."***

Reconciling speed with confidence

Claims like "faster" and "more trustworthy" tend to get framed as competing goals, but the tension is better understood as a routing problem rather than an inherent trade-off.

As decision-making becomes more autonomous, the rigor behind the verification supporting it has to increase in step. This is because there is no longer a human checkpoint to catch drift before it compounds.

"What we believe is the right solution is an intelligent network providing an intelligent way of routing this information with the right balance of speed and confidence," Gutting explains. "This independently sourced information is the one variable that needs to keep pace."



Verified data is the new trust layer beneath AI

As AI increasingly automates more of the decision layer across lending, audit and compliance, something still has to hold that layer up.^{7,13} Organizations will increasingly compete not just on how quickly they can decide, but on whether they can establish confidence in the information those decisions depend on.

“Verified data becomes the trust layer under the AI,” Gutting explains. ***“The more autonomous the decision-making, the more the weight of the verification has to be.”***

Verification is also becoming an ongoing practice as audit, lending and compliance workflows move toward automated review. Gutting suggests that “This doesn’t need to be an event, but more of a process.”

Verified data provides the trust layer that lets those automated workflows operate with greater confidence, while still leaving room for other controls and human judgment.

The industry is still early in this shift.

“We feel like we’re in the first inning of a lot of these AI models really being put into practice,” Smith says, suggesting organizations still have an opportunity to build verification into the architecture from the outset, rather than adding it later in response to a failure or regulatory requirement.^{7,13}

Conclusion

Verification: A required control

The financial industry’s trust problem has moved. The question is increasingly less about whether an artifact looks authentic and more about whether the underlying information can be traced to an authoritative source.

Documents are becoming weaker proxies for truth as AI makes convincing forgery easier, and detection alone will not solve that problem.^{1,2,3} It is a race where generation and detection evolve in

Verified data becomes the trust layer under the AI...

Jake Gutting,
Chief Product Officer,
Veriv



response to each other, so the finish line keeps moving. And once a decision is automated, an unverified input no longer causes just one mistake. Instead, it can be acted upon, propagated through other decisions and repeated before anyone has a chance to step in.

Together, these shifts point in one direction — verification needs to move upstream. Source-first verification provides the architecture for doing that, establishing trust before consequential decisions are made rather than trying to establish it after the fact.

Organizations have a choice right now, while building that gate still counts as a strategic edge. Wait too long, and the choice becomes a retrofit, done under regulatory or reputational pressure after something has already gone wrong.^{7, 13} Organizations that establish verification upstream can make it part of the decision infrastructure itself, rather than another control added after the workflow is already built.

So it is worth asking where, in your own workflow, a document rather than a source is still the thing being trusted. Where does an uploaded PDF, a self-reported figure or a customer-provided statement still stand in for the underlying record?

Veriv builds the infrastructure that establishes this kind of source-level trust, without asking requestors or responders to manage the complexity underneath it. The goal is simple — the organization making the decision should not have to chase down, interpret or manually validate the evidence it depends on. Rather, it should be able to request verified information and receive it directly into the workflow where the decision happens.

When the decision becomes automated, trust in the input cannot remain an assumption. It has to be established before the decision begins.

About Veriv

Veriv's unified financial verification infrastructure delivers truth, direct from the source - standardizing how financial data is requested, validated, and exchanged across audit verification, income and asset verification, and eligibility checks. Veriv is focused on bringing trust, consistency, and efficiency to one of the most critical functions in financial services.

For more information visit veriv.com.



→ Discover more



Sources

1. Financial Crimes Enforcement Network, "FinCEN Alert on Fraud Schemes Involving Deepfake Media Targeting Financial Institutions." FIN-2024-Alert004, Nov. 13, 2024.
www.fincen.gov/news/news-releases/fincen-issues-alert-fraud-schemes-involving-deepfake-media-targeting-financial
2. PYMNTS, "AI Cracked the Mortgage Verification System," July 10, 2026.
www.pymnts.com/news/artificial-intelligence/2026/ai-cracked-mortgage-verification-system
3. Regions Bank, "AI mortgage scams: How fraudsters target homebuyers and what they can fake now," June 3, 2026.
www.regions.com/insights/personal/article/ai-mortgage-scams
4. National Mortgage News, "Mortgage fraud case shows AI's fake-document risk," July 24, 2026.
www.nationalmortgagenews.com/news/mortgage-fraud-case-shows-ais-fake-document-risk
5. Laura Freeman, Feras Batarseh, D. Richard Kuhn, M S Raunak and Raghu N. Kacker, "The Path to Consensus on Artificial Intelligence Assurance," Computer, March 15, 2022.
<https://ieeexplore.ieee.org/document/9734265>
6. National Institute of Standards and Technology, "Reducing Risks Posed by Synthetic Content: An Overview of Technical Approaches to Digital Content Transparency" (NIST AI 100-4), Nov. 2024.
<https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-4.pdf>
7. Nupur Anand, Reuters, "U.S. bank regulators ramp up scrutiny of AI use at financial companies," June 12, 2026.
<https://www.reuters.com/business/finance/us-bank-regulators-ramp-up-scrutiny-ai-use-financial-companies-2026-06-12>
8. U.S. Government Accountability Office, "Fraud & Improper Payments."
<https://www.gao.gov/fraud-improper-payments>
9. Financial Industry Regulatory Authority, "Anti-Money Laundering, Fraud and Sanctions," 2026 FINRA Annual Regulatory Oversight Report.
<https://www.finra.org/rules-guidance/guidance/reports/2026-finra-annual-regulatory-oversight-report/aml>
10. Financial Industry Regulatory Authority, "Protecting Your Investment Accounts From GenAI Fraud," Jan. 15, 2025.
<https://www.finra.org/investors/insights/gen-ai-fraud-new-accounts-and-takeovers>
11. Australian Financial Review, "Team Australia must galvanise to stem \$4b in home loan fraud," June 28, 2026
<https://www.afr.com/companies/financial-services/team-australia-must-galvanise-to-stem-4b-in-home-loan-fraud-20260626-p60ahp>
12. Australian Financial Review, "CBA probes \$1b in suspected fraudulent home loans, calls in police," Feb. 27, 2026.
<https://www.afr.com/companies/financial-services/cba-probes-1b-in-suspected-fraudulent-home-loans-calls-in-police-20260223-p5o4mc>
13. Financial Stability Board, "Sound Practices for Responsible Adoption of Artificial Intelligence (AI): Consultation report," June 10, 2026.
<https://www.fsb.org/uploads/P100626.pdf>



**Any Verification.
Anywhere.**

veriv.com

