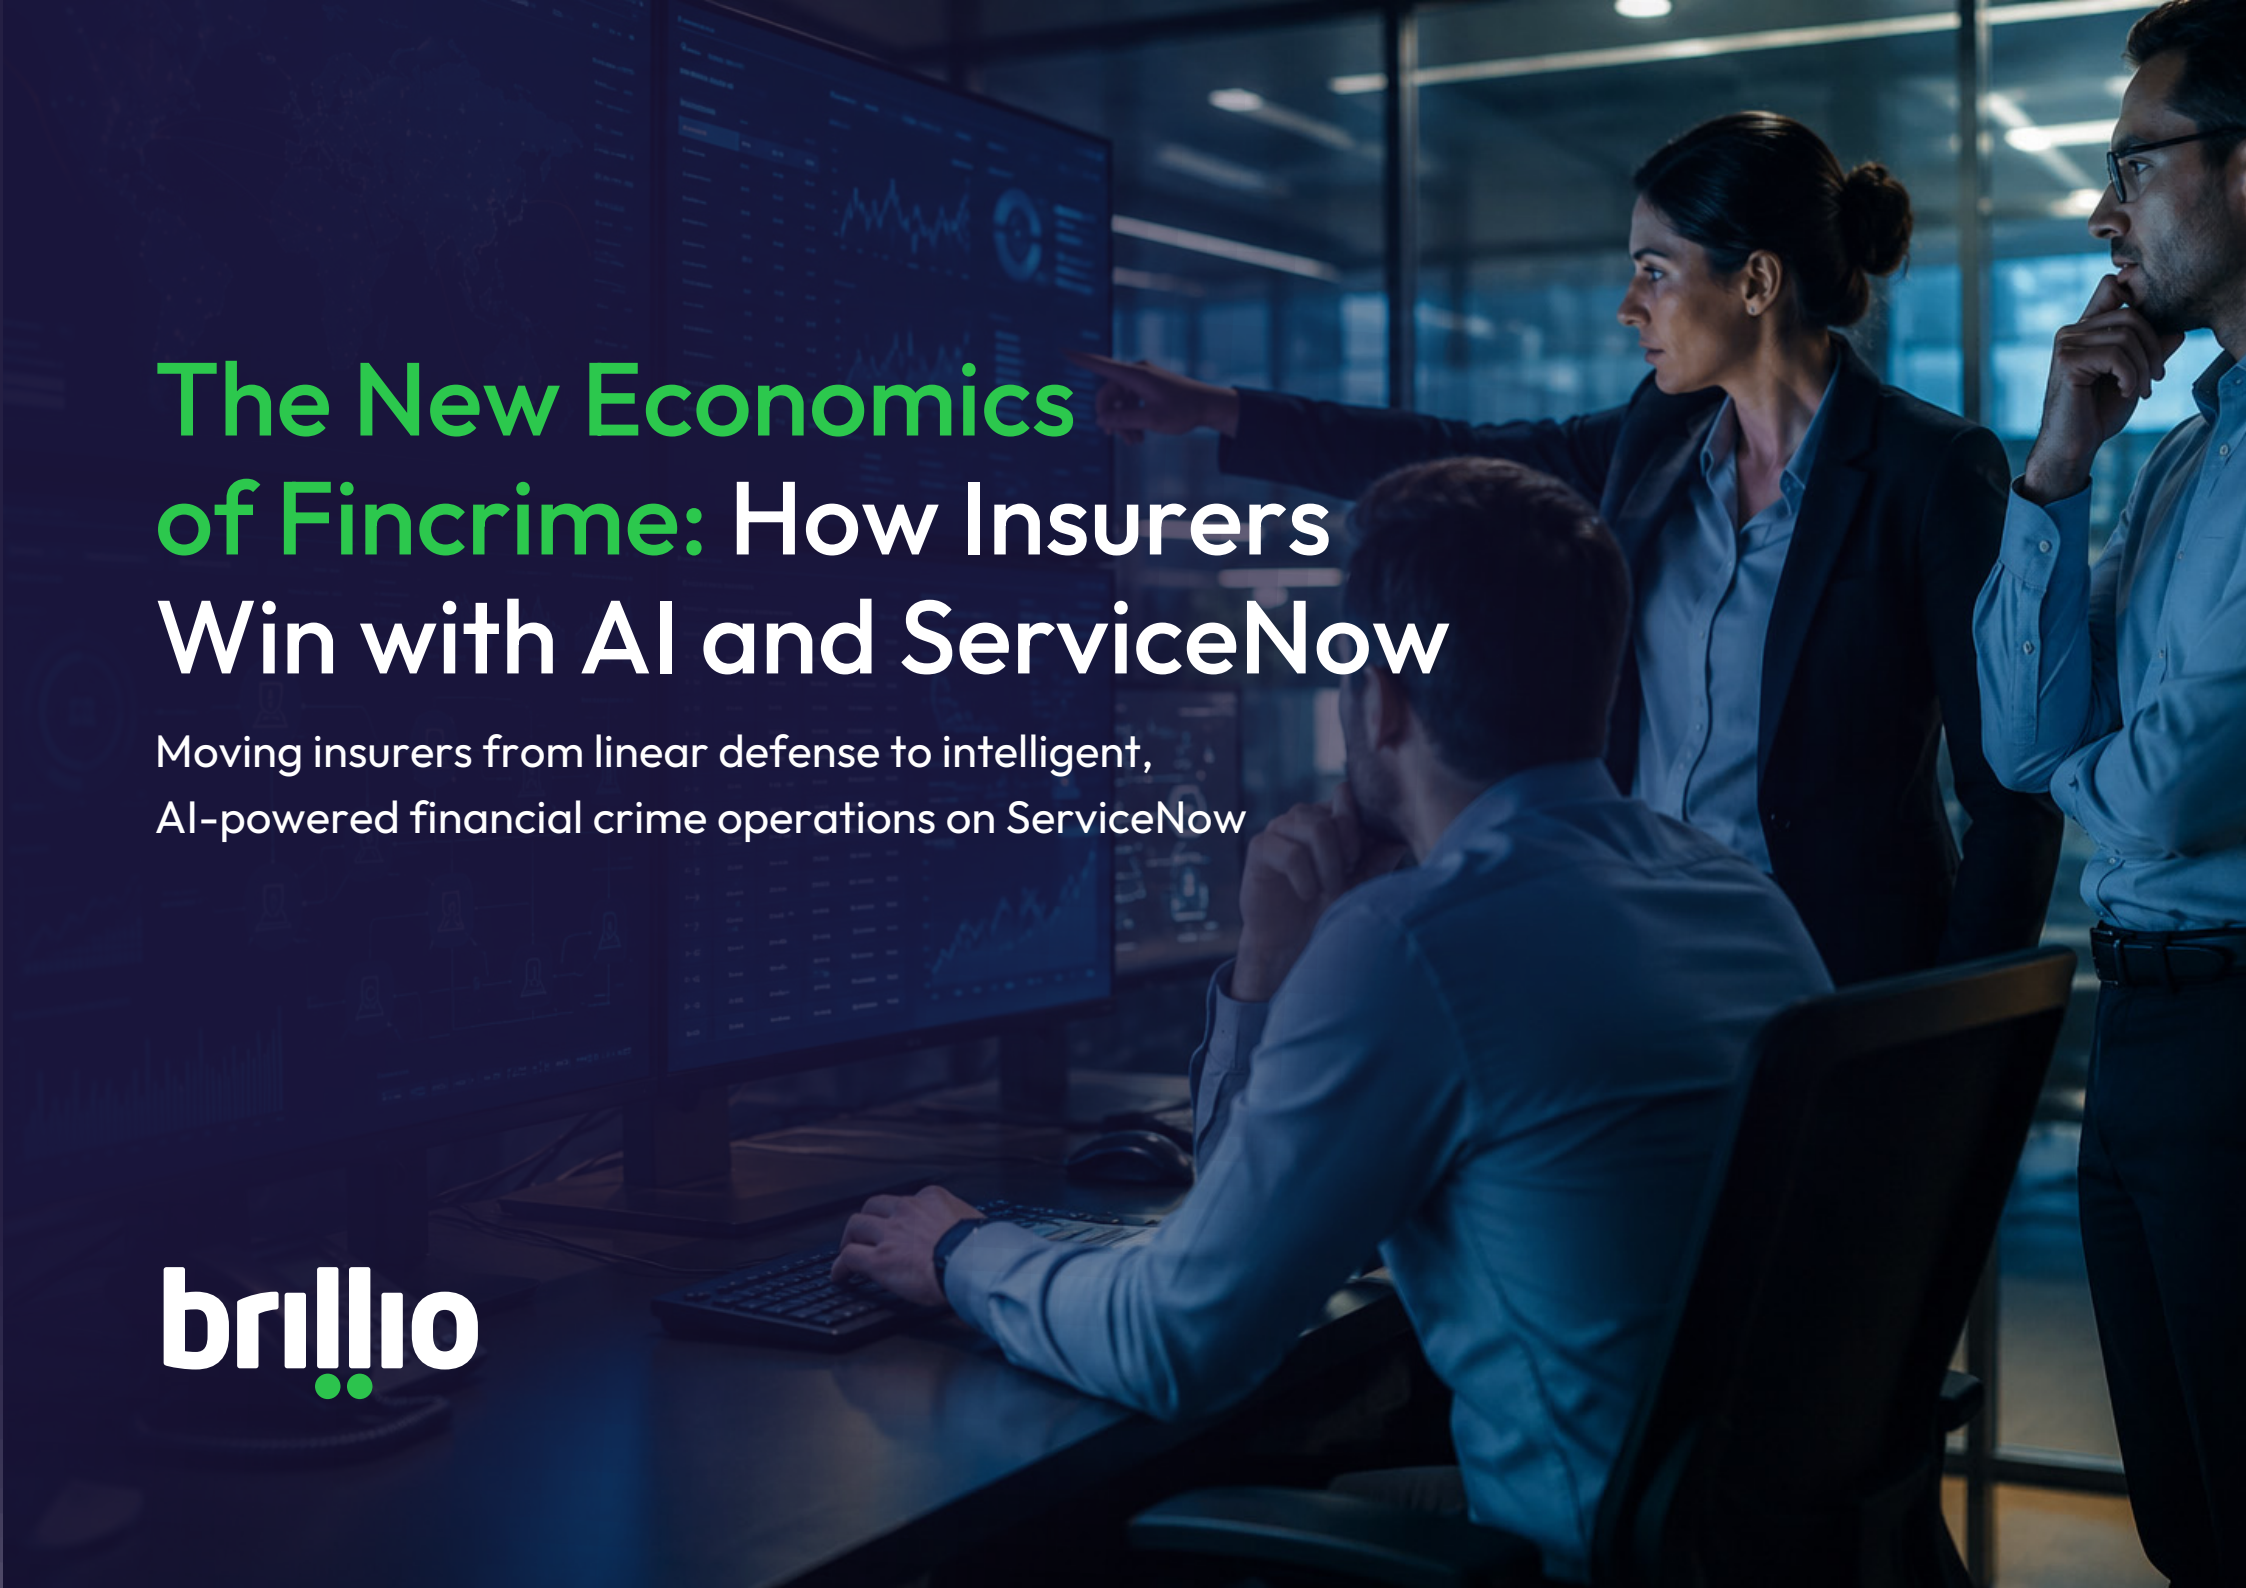

The New Economics of Fincrime: How Insurers Win with AI and ServiceNow

Moving insurers from linear defense to intelligent, AI-powered financial crime operations on ServiceNow

brillio

Fraud is scaling exponentially. Most insurers are still defending linearly.

For years, fincrime operations have been built on a simple equation: more cases, more analysts. That equation worked when fraud grew at a manageable pace, and threats followed predictable patterns. It doesn't work anymore.

The threat landscape has moved on. Fraud rings now operate at machine speed, using AI to generate synthetic identities, clone voices, fabricate documents, and exploit gaps between systems that were never designed to talk to each other. Defense, by contrast, still runs on spreadsheets, email threads, and the tribal knowledge of a few experienced investigators. The asymmetry is widening every quarter, and the cost of standing still is no longer hypothetical.

The challenge: a $30B problem getting worse

Data suggests fincrime is one of the few areas where every metric is moving in the wrong direction at the same time. Loss ratios are climbing. Investigation cycles are lengthening. Regulatory scrutiny is intensifying. And the analyst pool that traditionally absorbed all of this is harder to hire, harder to retain, and harder to scale.

The data backs up what we're hearing in those conversations. Across the industry, the numbers point to a defense model that is losing ground:

- 74% of insurers report fraud cases are steady or rising
- 1 in 30 claims is suspected to involve fraud
- Claim processing time triples when fraud is suspected, from around 21 days to 68
- $30B+ is lost annually to fraudulent P&C claims in the US
- €13B+ in annual fraud losses across the EU
- Only 23% of executives believe their compliance programs are highly effective

What makes this particularly difficult is that AI is now arming both sides of the fight. Fraudsters use it to industrialize attacks. Insurers, in most cases, are still figuring out where to apply it. The teams that close that gap first will operate at a structurally lower cost and risk profile than those that don't.

Why the current model is breaking

When we look inside fincrime operations, the same four bottlenecks show up almost everywhere. They aren't process problems that can be fixed with better training or tighter SLAs. They are structural, and they compound on each other.

The first is reactive threat detection. Investigations only begin after a lead has been manually surfaced, which means the team is always responding to something that has already happened rather than spotting it as it emerges.

The second is operational drag from fragmented data. Analysts spend a large portion of their day stitching together information on people, accounts, and addresses from systems that were never designed to share context, which means complex criminal networks often stay hidden in the gaps between silos.

The third is inconsistent risk prioritization. Case triage relies heavily on individual judgment, which leads to investigators working cases that should have been triaged out, while genuinely high-risk threats sit in queues waiting for attention.

The fourth is knowledge silos. Investigations depend on a handful of experts rather than a repeatable, data-driven playbook, and when those experts move on, the institutional intelligence leaves with them.

The combined effect shows up on the P&L as rising operational costs, on the risk register as delayed response times, and in regulator conversations as growing exposure that's hard to defend.

Why ServiceNow, why now

ServiceNow already runs across IT, HR, GRC, and customer workflows in most large insurers. It's enterprise-grade, governed, audited, and trusted by the security and compliance functions that would otherwise have to vet a new vendor.

It also means no new vendor investment, since we use the platform our customers already own. Enterprise-grade security and governance come built in through ServiceNow's access model and upgrade discipline. The solution is scalable and extensible, with native integration into KYC, CRM, claims, and payments systems. And time to value is shorter, because teams are already familiar with the interface and the operating model.

The strategic value here is centralization. One platform, one system of record for every lead, case, entity, and outcome, with a single audit trail behind every decision.

Inside the solution: a unified fincrime operating model

Our solution brings four layers together in a single workspace, designed to mirror how fincrime teams actually work rather than how vendors typically structure their feature lists.

At the top, **industry applications** cover the visible workflow: detect and ingest leads, manage group investigations, monitor suspects, capture outcomes, and store supporting documents. Underneath that, the **customer engagement layer** handles omnichannel lead reporting across voice, web, email, mobile, chat, and social, with an integrated agent workspace for routing, assignment, and collaboration.

The **financial crime operations layer** is where most of the day-to-day investigative work happens. It provides structured workflows for cases, leads, suspects, alerts, entities, documents, and a map view that visualizes every connection between leads and cases end-to-end. Finally, **the future-proof intelligence layer** holds the AI agents that work continuously in the background: data quality, customer profile aggregation, behavioral risk scoring, watchlist screening, fuzzy name matching, screening hit disposition, and proactive alert triggering.

The result is that disjoint processes become standardized workflows, fragmented data stores become a relational financial data model, and disconnected teams become a coordinated service organization with managed accountability.

The AI agents doing the work

Underneath the workflow sits a network of purpose-built agents, each with a defined job and a clear handoff to the next.

A Monitoring Agent scans payment flows, transactions, KYC documents, and customer behavior in real time, using ML models like Isolation Forests and AutoEncoders to detect anomalies. A **Lead Triage Agent** evaluates each alert, enriches it with context, and creates structured incidents with human-readable summaries. An **Assessment Agent** uses RAG-based retrieval to pull similar historical cases and prioritize based on context. A **Dedup Agent** identifies duplicate leads before they multiply investigative effort. A **Decision Gate Agent** auto-classifies high-confidence cases and escalates the rest for human review. A **Case Creation Agent** turns a final decision into a structured case in ServiceNow, complete with risk level, historical references, and an LLM-generated investigator briefing.

Every step is auditable, every decision is traceable, and every action is enriched with the context an investigator needs to make the next call quickly.

What our customers are telling us

The outcomes we're seeing in live use tell a consistent story. The platform is reshaping how teams spend their day, where their attention goes, and how quickly they can move from a flagged lead to a defensible decision. The most measurable improvements have shown up in four areas:

- Investigator efficiency is up 10 to 15%, based on direct feedback from internal fraud team leads
- Lead logging time is down roughly 50%, from around 20 minutes to 10 minutes per lead, which adds up to about an hour saved per person each week
- AML Ops teams are saving around an hour a day per user, through quicker mailbox logging, ATM allocation, screen navigation, and outcome capture
- Search speed is significantly faster, particularly when retrieving and referencing older leads
- The platform has held 100% uptime since launch, with users describing it as responsive, reliable, and quick to load

Beyond the metrics, the feedback we hear most often is qualitative. Teams describe seamless multi-case navigation through a tabbed interface, a more intuitive visual layout, and growing confidence that the platform can flex into adjacent use cases like sanctions, KYC refresh, and people-function integration. Just as importantly, the move from local tools like Rebus to an enterprise-backed platform has improved reliability, auditability, and security posture in ways that matter to CISOs and regulators alike.

Who this is suited for

This solution is built for fincrime functions that have outgrown their current operating model and need to industrialize without starting over. If several of the following sound familiar, your team would benefit from this kind of transformation:

- You track fraud and financial crime across multiple systems, spreadsheets, or email threads
- Your case progression is inconsistent or slow, and you can't always explain why
- You suspect you're missing connections between people, claims, or policies
- Your teams would benefit from a single view of entities across the business
- You already run ServiceNow but lack a structured fraud investigation module
- You need stronger confidence that only the right people can access sensitive fraud data

The more of these that resonate, the more value there is in moving early.

Where this can go next

What we've described so far is a foundation, not a finished state. Once the operating model is in place and the data is flowing through a single system of record, the work becomes considerably easier to extend, and the next layer of value starts to appear quickly.

The most immediate evolution is AI-driven pattern recognition across entity relationships, which becomes possible the moment the data is unified. From there, cross-team fraud intelligence sharing inside the tool starts to break down the silos between fraud, compliance, and legal that have historically held insurers back. Automated assessments and next-best-action suggestions follow naturally, and so does deeper integration with KYC, claims, payments, and third-party screening tools.

The compounding effect matters here. Each capability added to the foundation makes the next one easier, faster, and cheaper to deliver. Over time, the fincrime function stops being a fixed cost that scales with case volume and starts behaving more like an intelligence capability that gets sharper with every investigation.

The fincrime threat landscape is exponential, and the defense models most insurers run today were never designed for it. We help insurers close that gap by unifying fincrime operations on ServiceNow and embedding AI exactly where it changes the economics of the work.

ABOUT BRILLIO

Brillio is The Enterprise AI Accelerator helping Fortune 1000 companies move from AI ambition to scaled impact, faster. Powered by our AI accelerator platform – Agentic Data and Application Management (ADAM), Brillio is one of the fastest-growing digital technology service providers, delivering transformation across five core workstreams: business-led transformation, customer experience transformation, AI and data engineering, digital engineering, and infrastructure engineering.

With 14 delivery locations across North America, Europe, and Asia and a team of over 6,000 customer-obsessed professionals, Brillio combines deep industry expertise, modern engineering, and accelerators to deliver measurable outcomes. Headquartered in Dallas, Texas, Brillio serves clients globally with a commitment to speed, scale, and measurable impact.

https://www.brillio.com/
Contact Us: info@brillio.com

brillio