

FINANCE



TOMORROW'S INSURANCE

Mastering Churn
Through Data

THE AI FATIGUE

What the institutions shipping AI
are doing differently

CONTENTS

3

The Year the AI Fatigue Set In

4

The Customer Left Three Months Ago. The Renewal Letter Just Hasn't Caught Up.

6

The Cost Center That Was Always a Revenue Surface

8

Can Your Data Actually Support an Agent?

10

Six Questions to Ask Before You Put an Agent in Front of a Regulator

12

Before and After: What a Finance Operating Model Actually Looks Like with AI in Production

14

The Words Are Doing Too Much Work

16

A Skeptic's Guide to Reading an AI Business Case

17

Choosing the Right Transformation Partner for a Multi-Cloud, AI-Enabled Future



The Year the AI Fatigue Set In

Ambition is not the problem. Pilots that never ship are.

Ask a bank COO or a head of transformation how they are feeling about AI in 2026, and the honest answer is tired. Tired of pilots that impress and do not ship. Tired of vendors promising that this platform, this copilot is the one that will finally change the operating model. Tired of explaining to the board why the productivity numbers from the pitch deck never showed up in the cost base.

Most of what has been called AI transformation so far has been AI tourism. A lot of movement, very little that stayed behind.

There is a reason. Institutions are caught between two broken approaches. Multi-year

consultancy programs that analyze every process before anything ships. And AI platform vendors promising that feeding a workflow into their engine will transform the business. One burns budget before value arrives. The other produces impressive demos and uncertain paths to production.

What is working looks different. Narrower scope. A working result in weeks. Discipline about the governance underneath. One workflow first, then the next.

That is the shift. Not a new technology. A new rhythm.

Welcome to the finance edition. No AI tourism.



The Customer Left Three Months Ago. The Renewal Letter Just Hasn't Caught Up.

Why insurance churn is a data problem before it is a pricing problem, and what the insurers actually fixing it are doing differently.

There is a conversation that happens in every insurance executive committee, usually in the second week of the quarter when the retention numbers come in. It goes something like this.

The retention team points at the discount budget and says the problem is that the quote went out uncompetitive. The pricing team points at the claims ratio and says the problem is that loss experience forced the rate. The digital team points at the NPS score and says the problem is that the servicing experience was friction-heavy. Someone in the room asks, not for the first time, whether the predictive model is actually predicting anything useful.

Everyone in the room is a little bit right. But the reason the conversation repeats every quarter is that everyone is looking at the wrong moment. By the time the renewal quote is on the table, the relationship has already been lost. The customer has had three unsatisfying service calls, two unanswered emails, a claim that took longer than it should have, and a portal that asked them to repeat information the insurer already had. The discount arrived months after the trust did not.

Churn is not a pricing failure. Churn is a data and servicing failure that shows up on the pricing line because that is the line that gets measured at renewal.

Why the signal was never in the dashboard.

Eighty to ninety percent of what an insurer knows about a policyholder is locked in unstructured form. Call transcripts, adjuster notes, chat logs, email threads, claims correspondence, portal behavior. Legacy systems could read the structured fields. They could not read the conversation. For two decades, the most informative data in the organization sat in systems that churn models were never designed to look at.

That is now changing, and it is the real story underneath the agentic hype. The ability to activate unstructured data at the scale of a policyholder book is what makes a real-time disengagement signal possible for the first time. Not a quarterly report on who left. A continuous read on who is about to.

The insurers making this work are doing three things at once. All three are necessary. None is a single vendor purchase.

One, they are unifying the data that already exists.

Not collecting more of it. The data has been there all along. The work is making it queryable and usable as a single view of the policyholder, across the structured systems of record and the unstructured signal buried in servicing and claims interactions. This is the foundational work most insurers thought they had done with a data lake project five years ago and have now discovered was not enough.

Two, they are letting deterministic and agentic logic share the same workflow.

Regulated steps run on deterministic rails, the kind audit and compliance already understand. Judgment-heavy steps, where context matters and rules do not cover every case, use reasoning over the unstructured signal. An orchestration layer decides which runs when, keeps the audit trail, and pauses for a human where a human decision is legally required. Not one model doing everything. Several kinds of intelligence doing what each is best at.

Three, they are rebuilding the retention team around the new signal.

The save desk that waits for the cancellation call becomes a proactive relationship team that intervenes weeks earlier, with context the old desk never had. The metrics change. The incentives change. The skills change. This is the hardest part, and the part most programs underinvest in, because it is not a technology purchase.

What a realistic first step looks like.

Nobody getting results has started by rebuilding retention end to end. They start with one segment and one churn driver. Usually the one where the predictive signal is clearest and the save rate is weakest. They build the unified view for that segment first. They run the new workflow for that segment first. They measure the retention lift against the control. And they decide, in weeks not quarters, whether to expand.

This is the practitioner work that sits between a platform capability and a business result. It is also where the partner choice actually matters.

The test three questions long.

Can you see the early signals of disengagement before the renewal window opens? Can you act on them inside the same week, not the same quarter? And when you act, can you prove to a regulator how the decision was made?

If the answer to any of those is no, the churn number on the board report next quarter is a symptom. The work starts further upstream.

Churn is not a pricing failure. Churn is a data and servicing failure that shows up on the pricing line because that is the line that gets measured at renewal.



The Cost Center That Was Always a Revenue Surface

Every retail bank's contact center is the most underestimated revenue surface in the building. Here is why the numbers keep missing that, and what the banks actually fixing it are doing.

There is a line in every retail bank's annual report about operating efficiency in the contact center. The number is always going down. Calls are shorter. Cost per contact is lower. Digital deflection is up. Everyone agrees the trend is positive. And then the retention team reports separately that deposits are leaving faster than new money is arriving, high-value relationships are churning to alternative providers, and the loyalty numbers are sliding in the segments that actually make the bank money.

These two reports are describing the same problem. Nobody quite says it.

The problem is that the contact center has been treated as a cost center for so long that every

investment decision in it optimizes for cost reduction. A call handled by a bot is a cost saved. A call handled in three minutes instead of seven is a win. A customer who hangs up and solves it themselves is the best outcome of all. The scorecard is built around hours, not relationships. And what it produces, predictably, is cheaper conversations that retain worse customers.

The banks getting this right have stopped thinking about the contact center as a line item and started thinking about it as the one place in the retail banking operating model where a customer who is about to leave can be talked into staying. Every call is a save opportunity or a cross-sell opportunity or a referral opportunity. The cost of missing those is larger than the cost of the call itself by an order of magnitude.

What the shift actually looks like.

Three things change at once, and none of them is a pure technology purchase.

The data gets unified for the moment, not for the dashboard.

Most banks already have a unified customer view. It updates overnight, feeds dashboards nobody reads, and is invisible to the agent on the call. The version that matters is the real-time read of the customer at the moment the call connects. What products they hold. What their tenure and lifetime value actually is. What changed in their account behavior in the last two weeks. What they called about last time. What the churn model says about them right now. The agent sees it before the second sentence of the conversation. The difference in what happens next is substantial.

The AI does the listening, not just the talking.

The first wave of contact center AI was all about generating answers and deflecting calls. The more valuable version is about understanding. Real-time transcription and reasoning over call content surfaces the intent behind the question, the frustration in the tone, the moment the customer mentions a competitor. The agent gets guidance they can use. The supervisor sees quality issues before the week-end report. The bank learns from every call in a way quality-assurance sampling could never match.

The voice channel stops being a legacy problem.

This is the part most bank CTOs are wrestling with in 2026, because the shift from adapter-based telephony to native voice inside the CRM is no longer optional. It is a forcing function. Every bank is going to make this change in the next 18 to 24 months. The banks treating it as a technology swap will pay for it twice: once to migrate, and once again later to rebuild it around AI. The banks treating it as the moment to redesign the contact

center around the new operating model will pay for it once and end up with the architecture that makes everything downstream actually work.

The honest read.

If your cost-per-contact is going down and your high-value deposit retention is going down at the same time, the AI investment is not the thing that is broken. The scorecard is. The contact center is still being measured as a place that spends money, and every decision flowing from that scorecard is making the real problem worse.

The banks getting this right have not replaced their agents with bots. They have redesigned the human workflow around a new kind of customer moment, and the technology underneath is the thing that made the redesign possible, not the thing that replaced the people.

Three questions worth asking in the next operating review. Does your contact center scorecard include a retention metric? Does your agent see the customer's lifetime value before they answer the call? And when the voice channel finally moves off the old adapter model, will it land in the architecture that makes AI possible, or the one that makes it an add-on project two years from now?

Cheaper conversations that retain worse customers. That is what a scorecard built around hours, not relationships, actually produces.





Can Your Data Actually Support an Agent?

Most FINS AI programs do not fail in the model layer. They fail in the data layer underneath. Here is why.

The pattern is consistent across banking, insurance, and wealth. An AI program gets approved. A model is chosen. A pilot runs. It produces an impressive boardroom demo. Then the demo cannot be scaled into production, and nobody can quite explain why.

The reason is almost always the same. The data the agent needs to reason over is not where it needs to be, not fresh enough to act on, not consistent across the systems that run the business, or not governed in a way a regulator will accept. The model was never the problem. The foundation was.

FINS is particularly exposed because the data lives in places that were not built for real-time AI access. Core banking platforms from the 1990s. Policy admin systems with overnight batch jobs.

Claims correspondence sitting in shared drives. Client records duplicated across eight systems because the book of business grew through acquisitions. None of that is a scandal. All of it is the reality the agent has to operate inside.

There is a useful way to think about what "ready" actually means, and it is not a question about the model. It is a question about four things underneath.

Freshness.

Is the data current enough to act on, not just current enough to report on? Overnight batch is fine for management reporting. It is not fine for an agent responding to a customer in the next 30 seconds.

Meaning.

Can your systems speak business language, not system language? A field called STAT_01 works for a developer. It does not work for a reasoning model. A semantic layer is the thing most institutions have never built, and the thing that quietly decides whether the agent works.

Unity.

Do you have one version of the customer, or five? Unified profiles are the oldest unsolved problem in FINS. Most institutions have several overlapping versions with reconciliation jobs running overnight. Every time the agent asks who this customer is, it gets a different answer depending on which system it asked.

Context.

Are your documents machine-readable, or are they PDFs in folders? Claims correspondence, underwriting notes, advisory documentation. If the agent cannot read them at runtime, the judgment that used to live in those documents cannot be automated without losing what made the judgment good.

None of these is a technology purchase. All of them are practitioner work. Most institutions are not there yet, which is why the pilots keep producing demos and the demos keep not becoming production.

The question to ask is not which model to buy next. It is which of these four you are furthest behind on, and whether your current program plan includes the work to close it.

Want the full diagnostic? Scan the QR code to take the FINS Data Readiness Assessment. Fifteen questions across five dimensions, with a maturity score and next steps tailored to your institution.





Six Questions to Ask Before You Put an Agent in Front of a Regulator

The best risk officers in financial services are not asking whether agentic AI is ready. They are asking whether the institution deploying it is. Here is the list they are using.

There is a meeting that happens in every financial institution right now, somewhere between the board's approval of the AI budget and the first production agent. It is the meeting where the risk officer, the compliance head, and internal audit sit across the table from the program sponsor and ask the question that stops most agentic initiatives. How do we know what this thing is actually going to do? It is a reasonable question. It is also a question that has a real answer, and the institutions getting AI into production are the ones that have built the answer before the meeting, not during it.

1 Are the leaders actually aligned on what the customer outcome is supposed to be?

This sounds soft and it is the single most common reason agentic initiatives fail in financial services. If the sponsor thinks the goal is cost reduction, the risk officer thinks the goal is compliance uplift, and the customer experience team thinks the goal is retention, the agent will optimize for something nobody asked for and everybody will be unhappy with what ships.

2 Do we understand what the platform can actually do, and where its limits are?

Not what the vendor said. What it actually does, in your environment, with your data, under your compliance posture. This is where the distinction between deterministic workflow logic and agentic reasoning becomes operational, not theoretical. Knowing which kind of intelligence runs at each step is the difference between an agent you can explain in court and one you cannot.

3 Will the people who used to do this work adopt the new operating model?

Financial services runs on human judgment networks that took decades to build. An agentic workflow that bypasses those networks instead of incorporating them is not faster, it is a failed deployment waiting to happen. The question is whether the people who will have to live with it were in the room while it was designed.

4 Is our culture comfortable with the speed of the execution plan?

Some institutions can absorb a four-week build-and-decide cycle and call it progress. Others treat the same timeline as a risk event because their decision culture is calibrated for quarters, not weeks. Neither is wrong, but pretending the second kind of culture is the first kind is how agentic programs die in phase two.

5 Is our funding tied to outcomes or to platforms?

The CFO question the board never quite asks out loud. If the investment is structured as a platform commitment with a multi-year horizon, the program is funding a vendor. If it is structured as a milestone-based commitment where the next phase funds only when the prior phase produced evidence, the program is funding a change. These look similar on a slide and behave completely differently in practice.

6 Do we actually know what data we have, what data we need, and how the agent's reasoning will be audited?

This is the question every regulator will ask. The honest answer requires knowing the ownership, quality, and governance posture of every data source the agent touches, plus a traceable record of why the agent made the decision, not just what it did.

Six questions. None of them technical, all of them diagnostic. If the team proposing your next AI initiative cannot answer all six, the right move is not to kill the initiative. It is to answer the questions first, because they will be asked eventually, and they are easier to answer before the agent is in production than after.

A financial institution that has answered these six has built something worth more than any single agent. It has built a governance layer that makes every subsequent agent easier, faster, and safer to deploy. That is the compounding return on doing the trust work early. It is also the clearest signal, when evaluating partners, of whether they have done this before. The ones who have will want to talk about these questions on the first call.





Before and After: What a Finance Operating Model Actually Looks Like with AI in Production

Not a technology comparison. An operating model comparison. The shift the institutions getting results have made, and the shift the ones still stuck have not.

The most useful thing a financial services executive can do in 2026 is stop asking what technology to buy and start asking what their operating model is supposed to look like on the other side. The technology conversation is downstream of the operating model conversation, and skipping that is the single most common reason AI programs

produce pilots and no production. Here is what is actually changing, across six dimensions, in the institutions getting AI into real workflows with real outcomes. Read it as a diagnostic. For each row, ask honestly which column your organization sits in today.

DIMENSION	BEFORE	AFTER
How work gets prioritized	Annual transformation roadmap, agreed in advance, funded in phases, measured against milestones set before anyone knew what the work would require.	Rolling portfolio of short, outcome-bound workstreams. Each commits to a result in weeks. Each earns the next round of funding based on what the prior produced. The roadmap is still there, but it is a living document, not a contract with the past.
How data gets used	Data lakes and warehouses optimized for reporting. Structured fields queried for dashboards. Unstructured signal stored but unused because no system could read it.	Unified data foundation that makes structured and unstructured signal usable in the same workflow. The read of the customer is continuous, not quarterly. The signal buried in call transcripts and adjuster notes is now an input to real-time decisions.
How AI shows up in the workflow	A new tool on an agent's desktop. A recommendation in a dashboard. A chatbot bolted onto the portal. Humans doing the same work, with an assistant.	Several kinds of intelligence running in the same orchestrated workflow. Deterministic logic for regulated steps. Reasoning for judgment steps. Human review where a human decision is legally or ethically required. The workflow is designed around the capability, not retrofitted.
How governance works	Compliance reviews after the fact. Risk officers brought in at the end to approve what has already been built. Audit trails assembled retroactively from system logs.	Trust architecture designed into the workflow from day one. Guardrails defined before the first agent runs. Human-in-the-loop points specified as part of the design. Audit trails are a native output, not a reporting exercise.
How teams are organized	IT runs the technology. Business runs the process. A project team meets in the middle and hands things back and forth across a fence nobody admits is a fence.	Small cross-functional working groups own specific workflows end to end. Business, technology, risk, and domain experts in the same room, with decision rights, for the duration of the work. The fence is gone.
How success is measured	Project completion, on time, on budget, against the scope agreed at the start. Post-go-live value realization studies nobody reads.	Flow efficiency. Hours of work migrated from human to automated. Cycle time reduction on specific workflows. Revenue retained or captured that would otherwise have been lost.

Nobody is in the second column for all six rows yet. The institutions getting results are in the second column for two or three and are working their way across the rest. The institutions that are stuck are usually in the first column for all six, with a large AI budget pointed at the first row and an expectation

that the other five will follow on their own. They will not. The operating model has to shift together. Move first in the row where your current model is most clearly the thing holding the rest back. In most institutions, that is either the data row or the governance row.



The Words Are Doing Too Much Work

A plain-English decoder for the terms in every agentic AI pitch deck, and what to watch for when a vendor uses them.

Twelve words you will hear in every AI conversation in financial services this year. Here is what each one actually means, and what to watch for when someone is using it to move a deal forward.

Agentic AI

A system that can reason over a goal, take multiple steps to pursue it, and act on external systems without a human in the loop for every step.

Watch for: "Agentic" being used to describe tools that are actually copilots. If a human has to approve every action, it is a copilot, which is fine. It is just not agentic.

Deterministic AI

Rule-based, predictable logic where the same input always produces the same output. Not new, not glamorous, and the reason most regulated workflows work at all.

Watch for: Vendors treating deterministic as a bad word. In financial services, the regulated steps of any workflow should be deterministic by design, not by accident.

Large language model (LLM)

The reasoning engine inside most agentic systems. It generates language, interprets unstructured inputs, and makes judgment calls in context.

Watch for: The assumption that a better LLM solves a worse process. A better LLM applied to a broken process produces a faster broken process.

Retrieval-augmented generation (RAG)

A technique that lets an LLM answer questions using specific documents or data sources rather than its general training.

Watch for: *RAG described as a moat. It is not. The moat is the quality of the retrieved data, not the retrieval.*

Orchestration layer

The software that decides which step runs when, whether that step is deterministic logic, an agent call, a human review, or a legacy system action.

Watch for: *Products that have an impressive agent and no orchestration layer. They do not scale past the demo.*

Trust layer

The governance, guardrails, audit, and data-access controls that sit around an agent to prevent it from doing things it should not.

Watch for: *The word being used without specifics. A trust layer that cannot tell you why an agent made a decision is a marketing term, not a control.*

Human-in-the-loop

A workflow pattern where the agent does most of the work but pauses for a human to approve before the commitment is final.

Watch for: *Human-in-the-loop as an aspiration rather than a design. If the workflow was not built with the pause points in place, adding them later is a rebuild.*

Guardrails

Hard rules that prevent an agent from taking certain actions, like refunding above a threshold without approval or accessing data outside its scope.

Watch for: *Guardrails described in general terms. The useful version is specific: what this agent cannot do, and what happens when it tries.*

Unified customer view

A data platform that brings structured and unstructured customer data into a single queryable layer. What is new is that it is now a realistic foundation for real-time AI action, not just a reporting surface.

Watch for: *A unified view that covers only the easy structured fields. The value is in the unstructured signal most systems have never read.*

Zero-copy

A data access pattern where one system can use data from another without moving or duplicating it. Reduces the hidden integration cost of most AI programs.

Watch for: *Zero-copy as a license feature rather than an architecture commitment. The difference matters.*

Fusion team

A cross-functional working group of business, technology, and domain experts who co-design a workflow. The successor to the traditional IT project team.

Watch for: *Fusion team as a rebrand of the existing project team. If the same people are in the room with the same decision rights, nothing has changed.*

Proof of Business

A short, hands-on engagement that produces a working prototype and a quantified business case, not a slide deck. Used as a decision point, not a destination.

Watch for: *Proof of value framed as a pilot. A pilot tests feasibility. A Proof of Business tests whether the investment case holds up. These are different things.*

A Skeptic's Guide to Reading an AI Business Case

The numbers in the vendor deck are real. They are also almost never yours.

Every AI pitch deck in financial services contains a slide with a number on it. Forty percent cost reduction. Fifty percent productivity uplift. Three times faster cycle time. The numbers are usually sourced from the best outcome from the best engagement in the best conditions the vendor has on record. Reading them as a forecast for your own institution is how budgets get approved for results that never arrive.

The realistic ranges for financial institutions that have genuinely moved AI into production look more like this.

Cost-to-serve reduction in the contact center: ten to thirty percent over 12 to 18 months. The low end is what you see when the technology is deployed but the operating model did not change. The high end is what you see when the human workflow was redesigned around the new capability. The difference is change management, not technology choice.

First Contact Resolution and handling time: five to fifteen points of FCR uplift, fifteen to thirty-five percent reduction in handling time, on the specific

use cases the AI is actually running. Ask what percentage of total case volume the quoted number applies to. The honest answer is often five to ten percent.

Churn and retention in retail banking and insurance: two to eight percentage points of retention lift on the segments where the predictive signal is strongest. Nobody has moved the whole book by eight points.

Advisor and agent productivity: twenty to forty percent time reduction on the tasks the AI is handling. The trap is mistaking task-level productivity for role-level productivity. Build every business case on role-level math.

Three questions for every AI business case. Where does the number come from, specifically? What had to be true about the operating model for the result to land? And if we got half of it, in twice the time, on a narrower slice of the business, would the investment still be worth making?

If the answer to the third is no, the investment is being made on hope.





Choosing the Right Transformation Partner for a Multi-Cloud, AI-Enabled Future

Most enterprise transformation programs do not stall because the technology fails. They stall because the operating model around the technology was never designed to hold it. Data sits in places where it cannot be activated. Integration layers are fragile. Governance is unclear. And when AI capabilities are layered on top, the cracks widen rather than close. Data sits in places where it cannot be activated. Integration layers are fragile. Governance is unclear. And when AI capabilities are layered on top, the cracks widen rather than close.

Choosing a partner used to be a question of which firm could implement a specific cloud. That question is no longer the right one. The decision now is which partner can help you make multiple clouds work as a single operating model, with data, workflows, and decisions that move cleanly across the business.

What to look for: Operational readiness, not feature checklists

Building an Agentic Enterprise requires more than enabling AI features. It requires a partner who can integrate, govern, orchestrate, and scale without introducing instability. Enterprise-grade automation does not operate inside a single cloud. It moves across customer journeys, supply chains, revenue processes, service operations, and marketing ecosystems. If one layer is weak, performance degrades. If integration is fragmented, processes stall. If governance is missing, risk increases.

The right partner is the one who can tell you, before the work starts, where your data will hesitate, where your processes will stall, and where your governance will be tested. Automation delivers value only when it operates inside structured environments, and structured environments are the result of deliberate design, not optimism.


Proven execution across the clouds that run your business

Cross-cloud delivery is where most partner conversations get vague. A partner can claim breadth on a website. The harder question is whether the same teams have stood up each of those clouds in production environments in your industry.

This is the standard to which OSF Digital holds itself and where our credentials are relevant. We hold Expert-level Salesforce Competencies across

Commerce, Marketing, Sales, Service, Platform, and Data 360, alongside Expert recognition in Consumer Goods, Retail, Manufacturing, and High Tech, with additional Accredited Competencies in Agentforce, Analytics, Financial Services, and Media. That coverage matters because the teams designing your commerce experience understand how it connects to service. The people configuring your service operation understand the data your marketing and sales teams rely on. The architects integrating your systems understand the industry workflows underneath them.

Products



Industries



Experience that prevents failure before it compounds

Enterprise transformation involves interconnected decisions spanning systems, teams, data, and processes. The cost of getting those decisions wrong is rarely visible the moment they are made. It shows up later as the AI initiative that cannot find clean data, the customer journey that breaks at a handoff, or the new capability that cannot be governed at scale.

A partner with real depth across the architecture helps you avoid those costs before they compound. The right partner is not the one with the most awards on a website. It is the one whose people have done this work before, in your industry, across the clouds you actually use, and can tell you what will break before it breaks.

If you are evaluating partners for a multi-cloud or AI-enabled transformation, connect with OSF Digital to assess your current readiness and define a practical path forward.

About OSF Digital

OSF Digital is a Salesforce-first consulting partner helping organizations turn AI, data, and trust into measurable business impact. With 20 years of hands-on Salesforce experience and deep multi-cloud expertise, OSF guides clients from strategy through implementation to managed services across financial services, retail, manufacturing, and more. Known for award-winning delivery and a distinctive partnership mindset, OSF acts as an expert extension of its clients' teams. With regional hubs across North America, LATAM, EMEA, and APAC, OSF combines global reach with personal, high-touch delivery.

Stay current.

Stay current on FINS and AI topics for 2026.
Scan the QR code to subscribe to the
OSF Digital newsletter or visit **osf.digital**.



