Research paper

State of AI

SOFICO

SOFTWARE DRIVING MOBILITY

Contents

Introduction & research approach

Introduction & research approach

The automotive finance industry has always been a slower mover. We've seen it with digitalization, subscription models, and cloud migration. Now, we're seeing it again with AI adoption. While this report will show that there are some outliers, the industry as a whole is well behind the curve compared to large enterprises in other industries, especially those in financial services.

Why is this? As a software company at the heart of automotive finance, SOFICO had to find out. In the first quarter of 2026, we interviewed twelve automotive finance leaders from six leasing companies and five captives to compare where they stand today and where they're heading. (One leasing company is represented twice.) Combining insights from these participants with quantitative studies and official stats from third parties, we present a qualitative report about AI maturity in automotive finance with a strong statistical backing.

Before we dive into concrete examples, we'll explore the industry's context and current AI adoption. Then, we'll explore the challenges automotive finance is facing as well as the opportunities that lie ahead. We'll conclude with a roadmap to catch up with global changes.

Where AI stands

from enterprise to automotive finance

Where AI stands

from enterprise to automotive finance

The global picture

The AI revolution is happening faster than any revolution before. In 2025, 55% of large EU enterprises were already actively using AI, according to Eurostat. And globally, 76% of the large enterprises NVIDIA surveyed were using AI actively. In both reports, that's more than any other company size.

But adoption figures can be deceiving. While strategic ambition is high across the board, operational readiness is not. Many enterprises are still building the foundations that AI depends on: clean data, governance structures, integrated tooling, and the internal skills to make sense of it all. The gap between what companies say about AI and what they've actually deployed is still wide for most.

55%

of large EU enterprises were already actively using AI

76%

of large EU enterprises report that AI has increased revenue

Financial services: mostly ahead of the curve

Within the enterprise landscape, the financial services industry has moved faster than most. According to NVIDIA , 65% of financial services firms (of any size) are actively using AI, a figure that puts the sector just ahead of the global average.

The sector also leads in responsible AI maturity, meaning that frameworks for governance, ethics, and regulatory compliance are more developed here than in most other industries.

What firms are actually trying to achieve with AI reveals another pattern. Operational efficiency tops the list for 42% of firms, followed by employee productivity at 33% and new business opportunities at 20%. Compared to the cross-industry average—where operational efficiency sits at 34%, employee productivity at 33% and new business opportunities at 24%—the instinct in financial services is clearly to reduce cost and effort before expanding capabilities. It's a more conservative sequencing, and one that, as we'll see, echoes strongly in automotive finance.

What makes automotive finance different

Automotive finance isn't simply another category within financial services. It carries a structural complexity that sets it apart. Regulated lending and leasing processes sit alongside asset-heavy operations with real residual value exposure. The business runs through dense partner ecosystems: OEMs, dealers, insurers, and remarketing networks; each with their own data, systems, and incentives.

On top of that, the financial services subsidiary in many organizations operates at arm's length from the group's central technology function. The tools, the data assets, and the strategic priorities don't always travel sideways. Just some examples from our interviews:

- a captive trying to catch up with an OEM program it had no part in building
- a leasing arm sitting outside a bank-wide AI rollout
- a regional unit yet to receive a single tool from its own group's AI team.

Automotive finance has traditionally waited for direction before acting. The instinct has been to let others absorb the risk of early adoption and move once the path is clearer. That caution is understandable, and it has usually worked out well. But this time they feel left behind and that might damage the business in the long term.

The instinct has been to let others absorb the risk of early adoption and move once the path is clearer. That caution is understandable but now the world is changing faster than ever.

The risk of adapting too slowly

It looks like the AI wave won't pause for sectors to catch up. Although a global challenger is still far off in many industries, incumbents are starting to be challenged by AI-native vertical integrators. In automotive finance, specifically, scale is no longer a barrier, so other tech companies can now more easily enter the market. Enterprises that once felt untouchable are slowly losing that sense of security.

We've already seen the impact of electrification on the automotive market with new brands like Tesla and BYD, who reached best-selling positions in just 20 years. Who says AI won't cause another revolution in the automotive industry? And who says it won't happen even faster this time?

OEM captives and leasing companies who want to thrive in the next decades have to act quickly. Early movers have already made foundational decisions: which data to structure, which processes to automate, which vendors to trust. These companies are now ready to scale. Once they get up to full speed, the gap with companies who are still exploring AI will widen exponentially.

The current phase of AI adoption in automotive finance

The current phase of AI adoption in automotive finance

Three phases of AI adoption

According to Deloitte, 34% of companies are already starting to use AI to deeply transform their businesses, create new products and services, reinvent core processes, or fundamentally change their business models. Another 30% are redesigning key processes around AI while keeping their business models intact. And the remaining 37% are using AI at a more surface level, with little or no change to existing processes. (Percentages add up to 101% due to rounding)

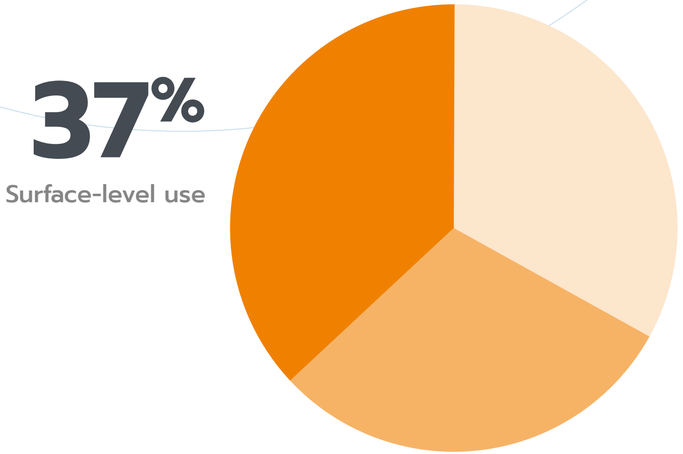

37%
Surface-level use

34%
Deep transformation

30%
Redesigning core processes

Wide maturity gap across the market

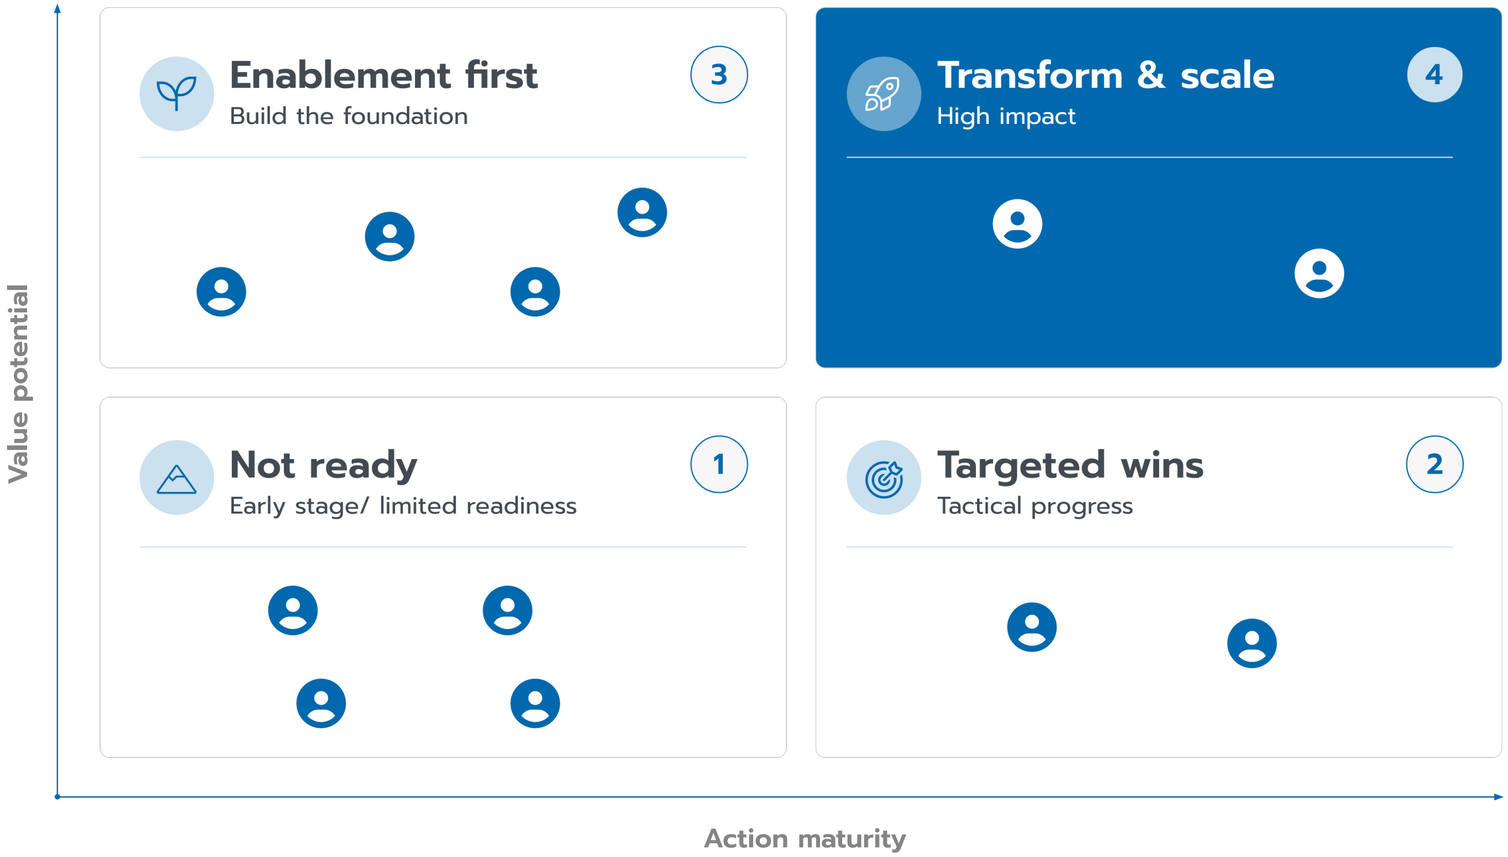

Our research showed that automotive finance, as a sector, sits predominantly in that last group. Many of our interviewees see their company's AI maturity as early stage: they are mainly exploring how to remove non-value-added work while they wait for better data and governance clarity. This confirms that automotive finance is behind other industries, but as we've said above, it's a rational starting point for an industry with high regulatory exposure and long technology cycles. Nonetheless, it means that the more transformative opportunities that could change the industry remain largely untouched.

UK LeaseCo

AI activity is largely limited to initial pilots and baby steps. We're not yet fully organized with end-to-end AI governance and processes.

EU OEM Captive

We have been on the trajectory already for 2-3 years. We have the platforms set up and we're now rolling out more and more PoCs in different arenas.

EU LeaseCo

We're early stage on AI and are still doing 'baby steps' with limited scaled deployment.

UK LeaseCo

Some technologies/use cases have been tested, but nothing is considered mature or scaled.

EU LeaseCo

We have used AI in a couple of cases but we are only now starting to build out more use cases.

UK LeaseCo

EU LeaseCo

We're starting the journey, with structured work beginning in 2025 via a task force generating AI use cases across the business.

EU OEM Captive

We already have AI use cases live in production, including OCR solutions and automated decisioning engines for financing and scoring. These are not pilots but operational systems

EU OEM Captive

Doing the same things, faster

The dominant pattern in AI adoption today is optimization. Across industries, most organizations are using AI to do what they already do, only faster or cheaper. Contracts are being reviewed more quickly. Emails are drafted in seconds. Reports can be generated without manual effort. The underlying processes stay the same, but AI is speeding up the process and trimming costs, albeit without significant impact.

Automotive finance is no exception. Based on our interviews, the industry skews heavily toward surface-level deployment with process-adjacent use cases rather than process-redefining ones. AI is being applied to tasks, not to the structures that organize these tasks.

What has already been deployed?

The deployments that show up most consistently across the players we spoke to are customer-facing chatbots and self-service workflows. These are the most visible AI investments, and in many cases the most mature. That's because they're relatively straightforward to deploy, the governance footprint is manageable, and the business case writes itself. They've been running long enough to have real operational data behind them.

Across our interviews, reducing back-office FTE was cited as a primary AI priority, with organizations consistently describing the goal as eliminating non-value-added tasks before anything else.

Below that, the picture becomes more varied. Fraud detection and machine learning for retention scoring are in use at a number of organizations, alongside OCR-based automation for document handling and back-office operations. These aren't new ideas, but the AI layer has made them faster and more scalable.

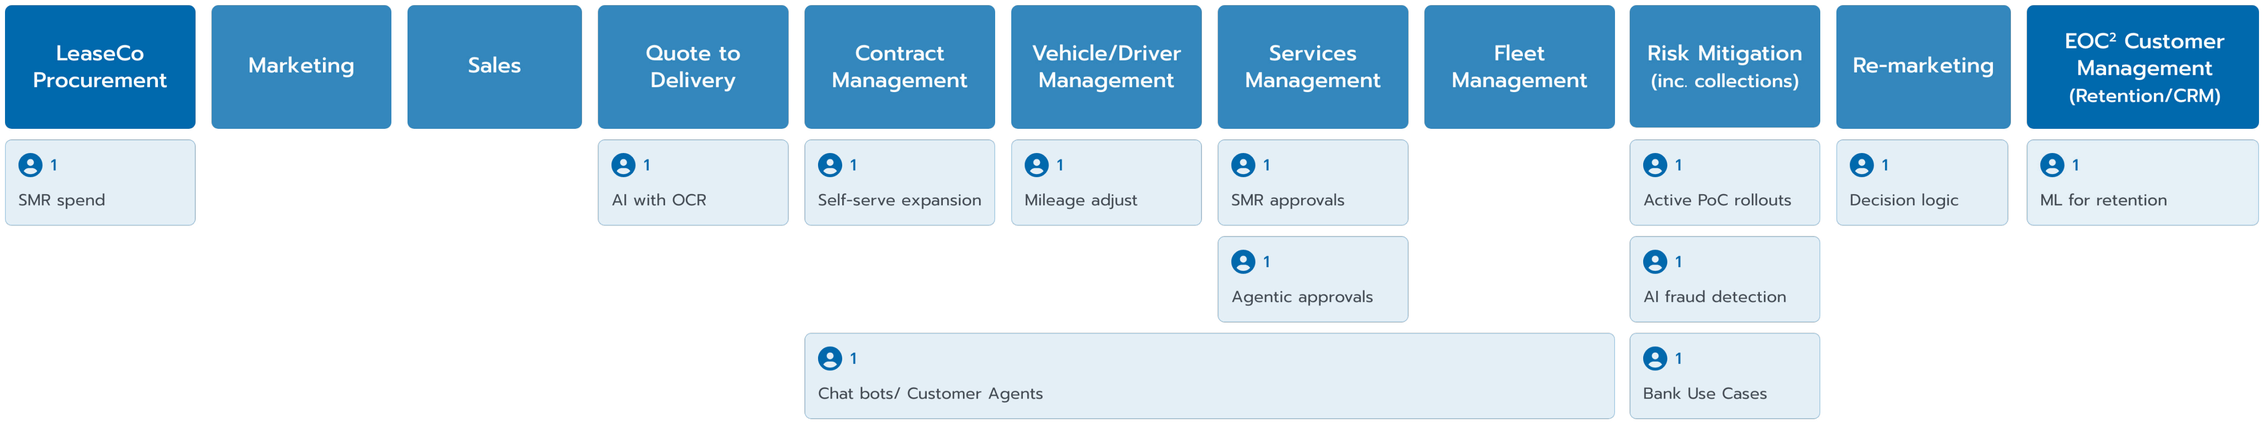

One captive we spoke to has built a full retention suite from historic contract behavior: a propensity-to-renew model that grades customers from low to high likelihood of renewal, a best-time-to-contact element that predicts when a customer typically acts rather than simply flagging end-of-contract proximity, and a next-best-car model that feeds dealer lead management with likely preferences before the conversation starts.

The same organization recently began integrating real-time usage data to extrapolate mileage trajectories earlier in the contract lifecycle. The commercial logic is straightforward: a customer is far more likely to accept a modest monthly adjustment at month eighteen than a significant settlement charge at month thirty-six. Moving from reactive end-of-contract identification to proactive in-life engagement is, in practice, a revenue protection tool as much as a customer experience one.

A fragmented landscape

The variation in AI maturity between players is striking. One leasing company already has forty or more dedicated AI staff, while other companies are limited to a Microsoft Copilot license for just over half of the employees, or worse, only a few employees who use ChatGPT informally. These organizations exist in the same industry, serving many of the same customers, and operating under the same regulatory constraints.

Crucially, this isn't just a matter of size or resources. Some smaller organizations have made deliberate investments and moved fast. Larger ones have sometimes been slower, caught between central group strategy and local execution.

In short, the automotive finance industry's AI adoption story is not one story. It's many stories, at very different stages of development. And that raises an obvious question: why are some automotive finance organizations struggling to keep up?

One leasing company already has forty or more dedicated AI staff, while other companies are limited to a Microsoft Copilot license for just over half of the employees

The structural barriers to further AI adoption

The structural barriers to further AI adoption

The pilot-to-production challenge

One pattern emerges repeatedly in our interviews, and it also showed in Deloitte's research: organizations have launched AI use cases without a clear path to measurable impact.

Even large, well-resourced organizations sometimes lack robust business cases for what they've built. The question of how it connects to revenue, cost reduction, or risk management turns out to be harder to answer than expected.

In other words, the gap between "we have AI" and "AI is generating value" is wider than reported maturity levels suggest. Many of the companies that would describe themselves as active AI adopters are, on closer inspection, running a portfolio of pilots that haven't yet translated into real operational impact.

What's making it so hard to go from pilot to production? Across industries implementation times have proven to be a difficult challenge due to the complexity that comes with scaling from an isolated environment to a company-wide role-out. The gap between expectation and reality for AI timelines is wider than most organizations anticipated. Deloitte's research puts numbers to a pattern that our interviewees described in almost identical terms: in mid-2025, only 25% of organizations had moved 40% or more of their AI experiments into production, yet 54% expected to reach that level within three to six months. Use cases estimated to take three months routinely stretch to a year or more when integration complexities surface. The optimism was consistent, but so was the delay in implementation.

The data readiness gap

Besides longer timelines, we noticed another pattern across our interviews: data quality and availability emerged as the single most recurring barrier to scaling AI in automotive finance.

> *The principal constraint is data architecture and access.*
>
> **OEM Captive**

Of course, there's a lot of data, but it's spread across systems that were never designed to talk to each other: dealer management systems, loan management platforms, OEM portals, remarketing databases, and so on. Each system only holds a piece of the puzzle.

On top of that, most of the data can't be even consulted by AI models without significant preparation. One company cites that "operationally, data is received as Excel files for a given day/month, requiring manual combination of multiple reports into a consolidated dataset." Especially OEM captives seem to be vulnerable as they are reliant on external parties, such as partner banks, for primary data.

> *Compared to the captives, we're in a much better position because there's a lot more requirements and scrutiny on the bank so the dataset is cleaner, the understanding of legal requirements is clearer, and there's more data on individual customers (especially when they are bank clients).*
>
> **EU Leasing company**

The delay between deployment and insights

Leasing and fleet management struggle with long contract lifecycles, which run two, three, four or even five years. That means the feedback loop for any AI system operating on that data is incredibly long. A system designed to optimize residual values or predict early terminations can't be validated quickly.

The result is that most players in automotive finance are not yet measuring impact in any rigorous way, and many projects haven't been running long enough to have results worth measuring. The question of whether their AI investments are working remains, for now, largely open.

Finding the talent to implement AI

The talent gap adds another layer of difficulty. Across every major study, insufficient skills consistently rank as one of the leading barriers to AI adoption.

- Eurostat found that 70% of companies that have considered using AI cite lack of relevant expertise as a reason for not proceeding.
- In NVIDIA's research, 38% of firms cited the lack of AI experts and data scientists as a prominent scaling challenge.
- McKinsey found that nearly 60% of respondents name knowledge and training gaps as the primary barrier to responsible AI implementation.

In our interviews there were no specific mentions of talent gaps, but there could be plenty of reasons for this. On the one hand, we could assume automotive finance companies don't face this challenge themselves as parent companies take care of AI hiring. On the other hand, many of the companies interviewed may simply not be hiring AI specialists yet as they've only recently started to plan for AI scaling.

70%
cites lack of expertise

38%
shortage of AI experts

60%
knowledge & training gaps

What separates scalers from laggards

Not everyone is facing the same pilot-to-production challenge. A smaller group of organizations has moved past the pilot stage and is generating measurable operational results. What they have in common is a clear data structure, dedicated AI teams, and meticulous prioritization frameworks. Here's a more detailed overview of the differences we found in our interviews.

What sets them apart?	Scalers	Laggards
Data infrastructure is treated as a prerequisite	**The most advanced organization we interviewed spent two years building a dedicated AI platform** with direct database access before attempting to deploy models at scale. Others framed their entire approach around the same principle: standardize processes and establish data readiness first, experiment with AI use cases second.	At several interviewed companies, data was still arriving as monthly extracts requiring manual consolidation before any analysis could begin. In others, AI roadmaps were blocked entirely behind infrastructure migrations that hadn't yet been completed.
Dedicated AI capacity	Advanced companies create **permanent AI functions** rather than running initiatives through shared IT resources. The most advanced players have structurally separated AI organizations with their own reporting lines and senior leadership accountability.	They rely on task forces, borrowed resources, or external consultants for individual projects. Institutional knowledge doesn't survive between pilots, and the internal muscle to move from proof of concept to production never gets built.
Structured use case frameworks with business value accountability	Prioritization frameworks are built before use cases are launched. These companies hold themselves accountable to the outcomes. **The most advanced player interviewed organizes every use case into one of four business-value buckets and tracks each proof of concept against that framework.** Others maintain formal backlogs with estimated value attached and report concrete results.	There's a long list of ideas, but no clear prioritization, no baselines, and no measurement. Without that accountability structure, it's almost impossible to build the internal business case needed to justify scaling.
Governance and central ownership	AI development is driven centrally, **with an explicit belief that it only makes sense at an organizational level rather than market by market.** One company completed a dedicated governance project before any production deployments went live, treating it as infrastructure rather than overhead.	Several early-stage players noted that they were not yet supported by an established organizational model or defined processes for AI. Initiatives were scattered across teams without clear ownership. One participant was direct about the consequence of not having this in place: without central steering, use cases were duplicating across geographies and costs were compounding.

Governance, risk and the regulatory landscape

A significant share of the industry still operates on a legacy of security requirements, regulatory constraints, on-premise infrastructure, and long technology cycles. That creates immediate friction for AI deployment, which increasingly assumes cloud-native environments, API-first architectures, and real-time data portability.

Many of the most capable AI tools on the market simply can't be connected to on-prem systems without significant engineering work. And moving to the cloud, while increasingly necessary, carries its own set of complications: data sovereignty, vendor dependency, security exposure, and the need to re-certify compliance across a new infrastructure stack.

What makes this even harder is that automotive finance players often answer to many governance authorities simultaneously. The local business entity owns risk appetite and budget. The group, OEM, or parent bank sets tooling whitelists, vendor approval processes, and data handling standards. Legal and compliance interpret regulatory requirements and manage liability. IT and security define architecture requirements and access controls. Any AI initiative has to clear all of them, and they rarely move at the same pace.

In practice, this multi-layered governance structure shapes nearly every dimension of what AI can actually do. Firstly, it determines which tools are permitted—approved vendor lists are increasingly standard across the players we spoke to. On top of that, procurement cycles for new tools can often last months in an industry that changes weekly.

> *If a vendor runs technology not approved by the corporate parent blueprint, implementation becomes difficult or impossible without heavy risk exception processes.*
>
> **EU OEM Captive**

Secondly, sovereign AI has moved from an abstract concept to a concrete planning constraint. Deloitte found that 83% of companies view sovereign AI as at least moderately important to their strategic planning, with 43% rating it as very or extremely important. Two thirds express at least moderate concern about reliance on foreign-owned AI infrastructure.

> *AI governance has been treated as a prerequisitie rather than an afterthought*
>
> **EU OEM Captive**

Thirdly, governance determines where human oversight is non-negotiable. Under the EU AI Act , automated systems used for credit or risk decisions must include meaningful human review at key decision points. In automotive finance, where AI-assisted decisioning for credit applications, residual value assessments, and contract pricing are core use cases, this isn't a simple compliance question. It defines what can be built, how it gets approved, and how fast it can move.

The deepest governance challenge, however, is agentic AI. Security and risk concerns are already the top barrier to scaling agentic systems: McKinsey found that nearly two thirds of organizations cite them as the primary obstacle, well ahead of regulatory uncertainty or technical limitations. The same research found that 74% of respondents identify inaccuracy as a highly relevant risk, and 72% cite cybersecurity.

The challenge with scaling Agentic AI use cases

Despite governance challenges, deployment intentions are moving fast: Deloitte found that 74% of companies plan to deploy agentic AI within two years, while only 21% report having a mature governance model for autonomous agents.

That gap is particularly dangerous in automotive finance. Agentic use cases here touch regulated credit decisions, financial transactions, and customer contracts. Deploying autonomous systems into those workflows without mature governance means taking on risk no one can measure, price, or hedge.

74%

plan to deploy agentic AI within 2 years

20%

have a mature governance model

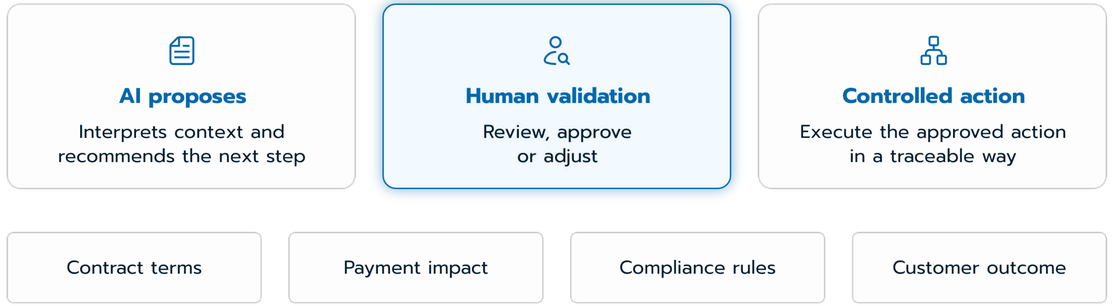

In conclusion, it's easy to see governance as a constraint, but the more useful framing is to consider it a design input. McKinsey's research shows that organizations assigning clear ownership for responsible AI exhibit the highest average maturity levels. The same research found that organizations treating AI trust as a core business capability rather than a compliance requirement are better positioned to scale AI to its full potential. Our interviews confirm it: the organizations with the clearest governance structures are also the ones moving fastest. Governance, done well, doesn't slow you down. It removes the ambiguity that slows you down.

Technology and vendor landscape

The AI tooling landscape is moving faster than most organizations can evaluate it. New platforms emerge, capabilities shift, and competitive positions that looked settled recently are already being contested. For automotive finance organizations operating within structured procurement cycles and multi-layered vendor approval processes, that pace creates a genuine problem. By the time a new tool has cleared the approval stack, the landscape it was assessed against has already moved on.

> “
>
> If CEOs see AI as a fad or something to avoid, they probably won’t survive
>
> **EU LeaseCo**

That friction compounds an existing frustration. Across our interviews, a consistent signal emerged: traditional core vendors have not kept pace with AI execution. The systems these organizations depend on for their most operationally critical data are not yet exposing it in the ways AI deployment requires. One interviewee was direct about the consequence: vendors perceived as lagging on AI execution and architectural readiness are already losing the confidence of their most advanced customers.

The fragmentation of the tooling landscape makes vendor selection harder still. There are many capable tools available but little integration between them. The risk of backing the wrong platform, or creating lock-in that limits flexibility as the landscape shifts, is pushing organizations toward caution and delay.

The preference for a “bring your own AI” approach is a rational response to all of this. Organizations want to own their model choices, their agent architecture, and their data flows. But that preference only works if the underlying platforms expose the right interfaces. When they don’t, organizations are left choosing between costly workarounds and a degree of vendor dependency they are trying to avoid.

Strategic opportunities for AI in automotive

Strategic opportunities for AI in automotive

Setting the frame: automation vs. AI

Not every intelligent system is AI and not every system needs to be built with an AI-layer. To avoid confusion, it's useful to separate AI from automation based on what each one actually does. Automation is deterministic: it follows rules, produces predictable outputs, and doesn't require model training. Given a specific input, it will always produce the same result. That makes it well suited to high-volume, structured, and repetitive tasks. It's lower risk, faster ROI, and far simpler to govern.

AI is non-deterministic. It recognizes patterns, makes predictions, and exercises judgment at scale. The output isn't predetermined but inferred from data. That capability is powerful, but it comes with meaningful trade-offs: more complex governance, slower validation, harder-to-explain decisions, and a much longer path from pilot to production.

> “Many requests are basic and should be solved via better data exposure and workflow automation rather than advanced AI. GenAI is most promising for complex, customer-specific ‘scheme queries’ where answers vary by customer, grade, and tenure.
>
> **UK LeaseCo**

The right question for any use case is therefore not whether AI is possible, but whether it's required. Many of the efficiency gains available to automotive finance organizations don't need a complex AI model. They need a well-designed rules engine and clean data. Reaching for AI when automation would do the job is a common and costly mistake, adding complexity and governance overhead where neither is necessary.

Not everything needs an agent

Responsible progress means choosing the right level of intelligence for the job

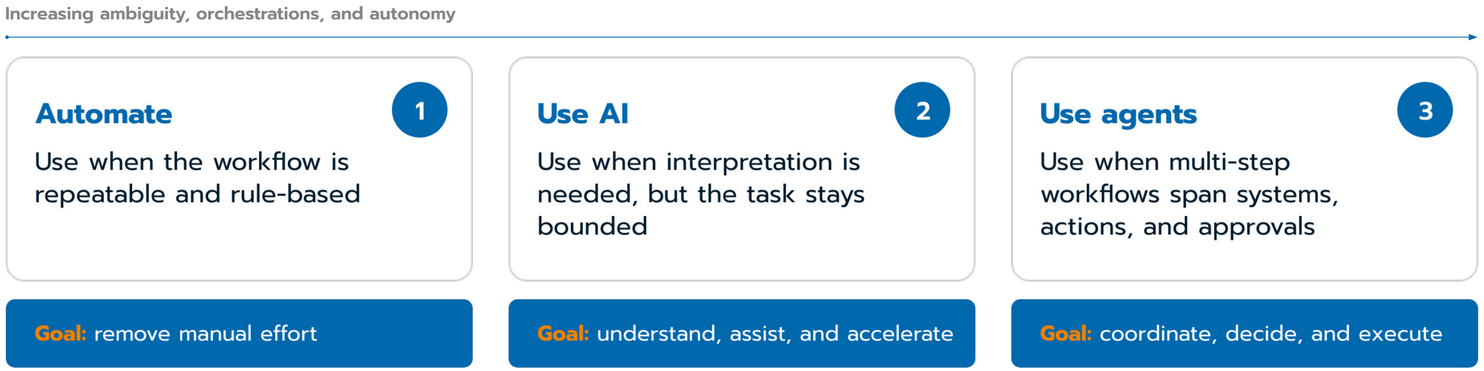

Having said this, let's zoom in on what companies are pursuing now and what other opportunities lie ahead. However, always keep in mind that rigorous measurement of results is still rare. Most players are not yet at the stage where they can report validated impact at scale. These are projections and signals; not proven outcomes.

What are most players doing now?

The most widely deployed AI capability in automotive finance today is also the most visible: automated customer service. Chatbots, FAQ deflection, and self-service workflows have become the default entry point for AI adoption across the industry.

A smaller number of organizations are moving into adjacent territory: compliance monitoring, exception flagging, and back-office process standardization, where volume and repetition make even modest automation economically significant.

These are valuable gains, but these use cases typically call for automations rather than complex AI systems. This means that many companies are less advanced in terms of AI than they think. Or they are overengineering some solutions and, hence, limiting the impact. Nonetheless, these use cases prove the value of intelligent tooling before committing to the infrastructure investment that more complex AI deployment requires.

The one exception is Copilot-style productivity tooling, which has become widespread quickly, often entering through group-level Microsoft licensing rather than any deliberate AI strategy. Its AI use case is real but mostly limited to making existing tasks faster and more efficient.

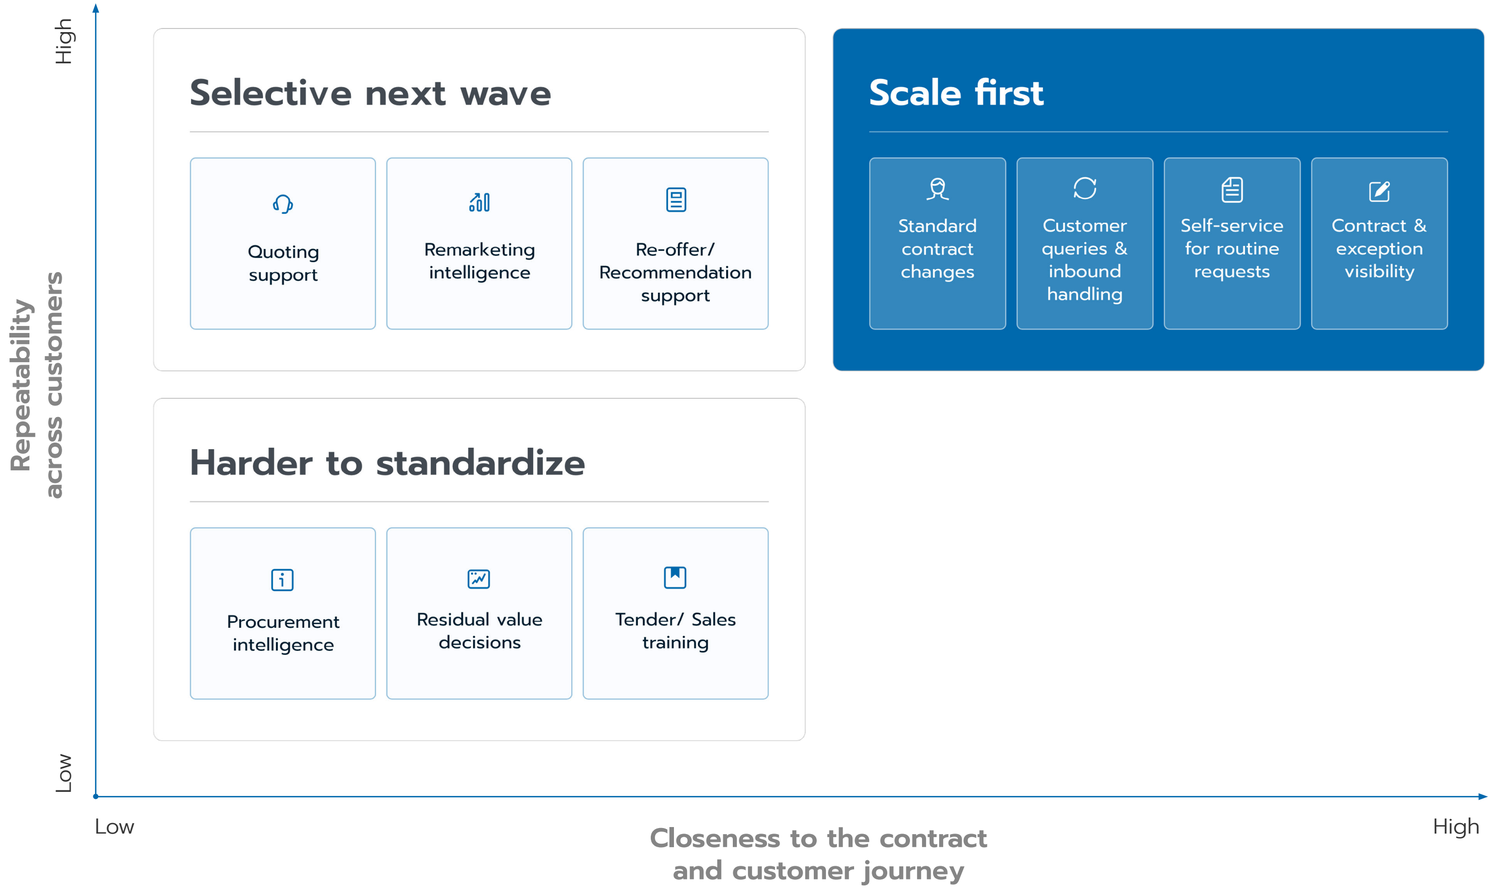

The next layer that's available near term

A second tier of use cases is already in active deployment at the more advanced players or in rollout at organizations moving faster than the peer group. What defines this tier is a step up in data requirements while most use cases are still limited to rule-based automation.

OCR and document processing is perhaps the clearest example. Automotive finance is contract-intensive by nature, and a significant share of those contracts, invoices, and damage reports still contain information locked inside unstructured documents. Extracting that information automatically—rather than through manual data entry—is a well-understood problem with well-tested solutions. Therefore, its efficiency potential compounds quickly at the volumes this industry operates at.

The same logic extends to basic operational tasks: approvals, status updates, and workflow routing that currently require human intervention but don't need human judgment. The rules are clear, and the variation is low, so removing the human from that loop doesn't reduce quality.

Machine learning for retention scoring is where genuine AI capability enters the picture. Predicting which customers are approaching end-of-contract and likely to churn, or identifying the right moment to start a renewal conversation, requires learning from historical patterns rather than following a fixed rule. A handful of organizations are doing this with real operational integration; others are at proof-of-concept stage.

And finally, there's vehicle operations, which is perhaps the most underinvested near-term opportunity across the organizations we spoke to. SMR authorization, maintenance approvals, and damage assessment are high-volume decisions that currently require significant human time per case. The data signals are structured enough for AI to assist, and the efficiency gains at scale are potentially substantial.

In practice:

UBench takes the manual work out of claim intake

At a fleet of 50,000 vehicles, claim intake means roughly 40,000 documents a year. Each requiring manual review, linking, and routing before any actual claims work can begin. Our research flagged vehicle operations as one of the most underinvested near-term opportunities in the industry: high volume, expert-dependent, and largely untouched by AI.

UBench, part of the SOFICO ecosystem, is addressing this directly. Its Casper capability operates inside the intake workflow, classifying and routing documents automatically, linking them to existing claims or creating new ones, and escalating only genuine exceptions to human reviewers.

Some early promising results: 92% of documents handled automatically, error rates reduced from 37% to 8%, and approximately €218,000 in annual savings estimated for a fleet of 50,000 vehicles.

> “The objective was not AI for its own sake, but removing manual coordination work and restoring operational flow across the claims lifecycle.
>
> **Jonas Seuntjens – CPTO, UBench**

- Fleet of 50,000 vehicles
- 40,000 documents per year
- Each requiring manual review, linking & routing

Document arrives
Claim intake

↓

Casper classifies
& routes automatically

↓

Linked to claim
Existing or new

↓

Human review
Exceptions only

92%
documents handled automatically

37% → 8%
error rate reduction

20%
estimated annual savings

One big underexplored near-term possibility

Regulatory and compliance reporting stands out in our research as a clear near-term opportunity that the industry has largely not yet acted on. The reporting burden in automotive finance is substantial and growing. Regulatory requirements vary by market, change frequently, and generate significant manual effort in data gathering, formatting, and submission. It's exactly the kind of problem where intelligent automation can reduce cost and error rates without requiring the organization to take on meaningful AI risk. And yet it rarely appears in the AI backlogs that automotive finance teams are building.

What advanced organizations expect from their vendors

The use cases in the first phases are largely deliverable through existing vendor capability and standard automation tooling. What follows is different. Competitive differentiation requires proprietary data, in-house model ownership, and a vendor ecosystem that enables rather than replaces internal capability.

The most mature organizations have stopped treating tooling providers as suppliers and started co-designing solutions with them. The vendor relationship at this level is collaborative rather than transactional, with both sides investing in making the architecture work. One leasing company we spoke to has already structured its vendor relationships by capability — one partner for telephony, another for customer interactions, a third for asset-related activities—and is actively scanning for the next gap to fill rather than waiting for a single vendor to cover everything.

It's a clear example of the overall signal from our interviews: organizations at this level want their platforms to behave as reliable data providers while they take care of the broader agent layer themselves. So, they want stable, well-documented APIs (and RAG) to expose contract facts, calculations, and workflow state. What sits on top of that, such as the models, the agents, and the orchestration logic are things they prefer to own in-house.

The main reasons are differentiation and control. Several interviewees expressed concern that if all players rely on the same platform for the same AI features, the competitive advantages they have built begin to erode. Back-office automation is fair game for vendor delivery, but anything touching tender strategy, commercial pricing, or customer retention is treated as proprietary territory.

Additionally, when an organization builds its own model on its own data, it has a much better understanding of the controls and exposures. Packaged AI features from a vendor do not always offer that transparency, which matters acutely in a regulated environment.

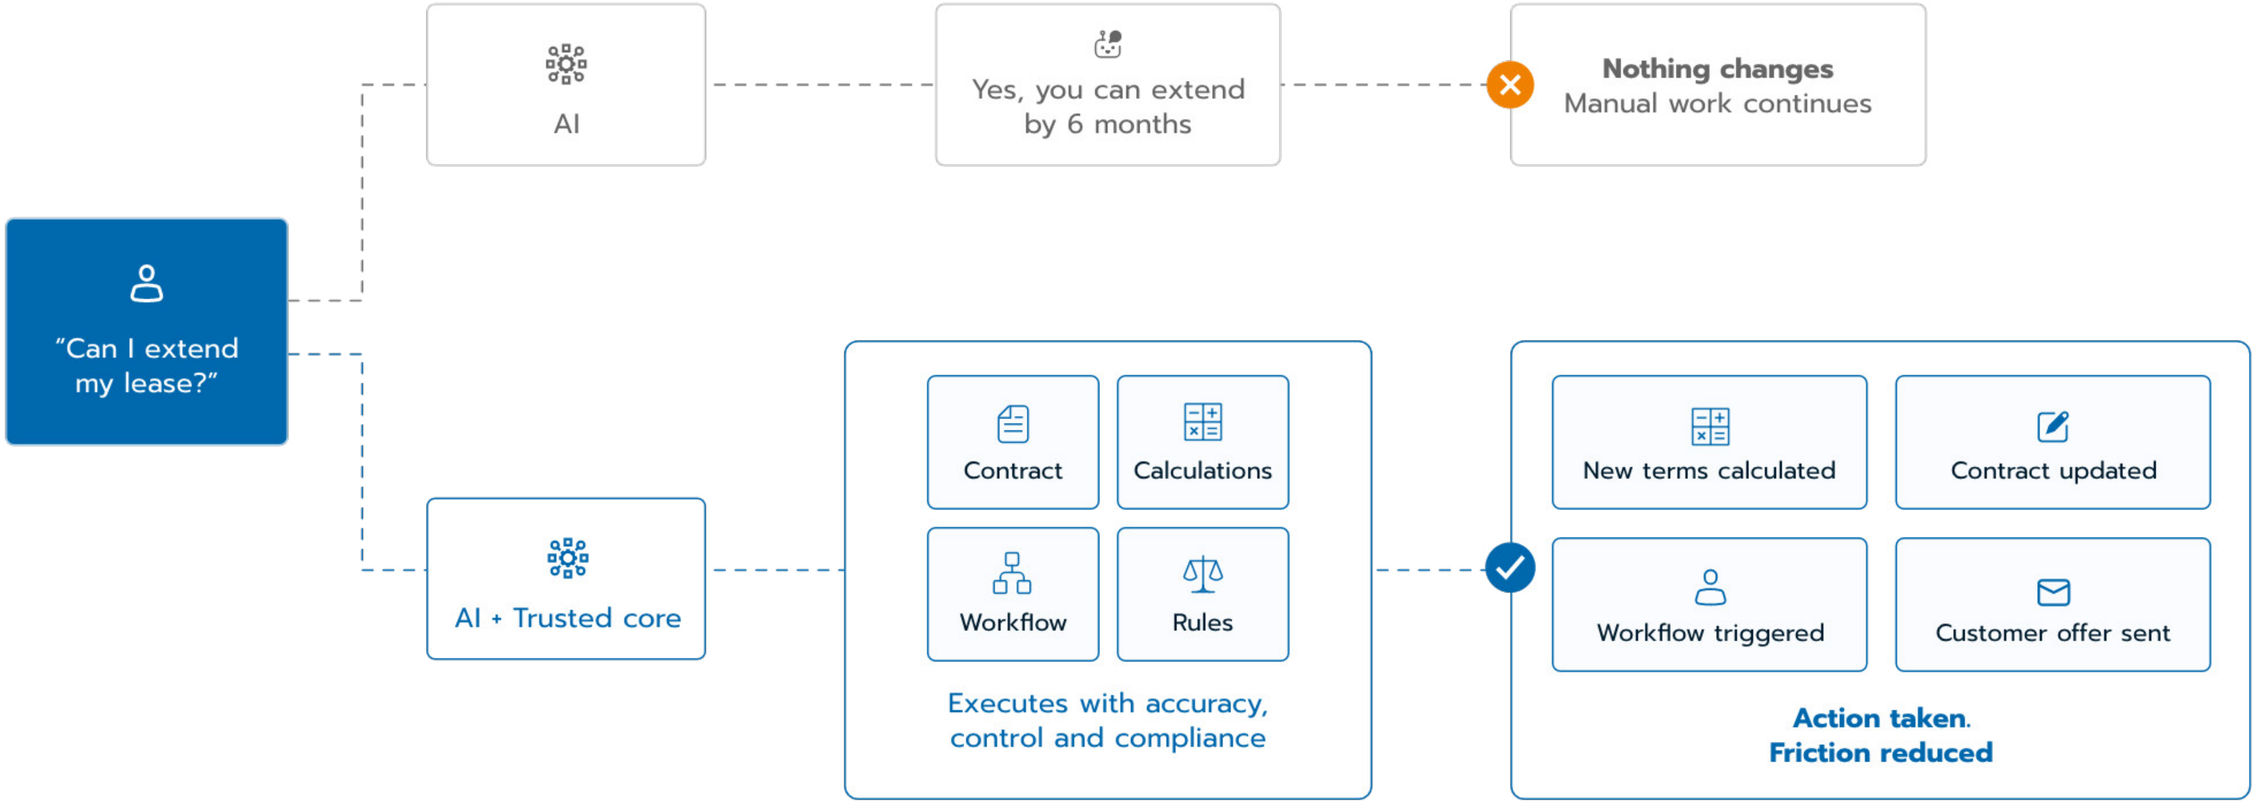

In summary, the leading AI integrators in automotive finance expect a robust, stable core that exposes the right data through reliable interfaces and stable, generic APIs. The rest is optional.

Competitive differentiation is a medium-term possibility

- **Intelligent remarketing and end-of-contract decisioning** are among the highest-cited opportunities in our interviews. Residual value management sits at the heart of leasing economics, and the decisions made in the final months of a contract, about whether to extend, refresh, remarket, or sell, have outsized impact on portfolio performance. Historically, those decisions have been made with incomplete data, under time pressure, and with significant variation across markets and individuals. An AI-model that can make the right decision for each vehicle, each customer, and each market condition in real time is a meaningful competitive advantage.
- **Next-vehicle recommendation and customer lifetime value modelling** operate on similar logic but face outward rather than inward. The question shifts from what to do with an expiring contract to what to offer that customer next, and when to make that offer. When managed effectively, the end-of-contract moment stops being a retention risk and becomes a commercial opportunity.
- **Fleet total cost of ownership profiling and portfolio health management** represent the kind of aggregated, cross-portfolio intelligence that has historically required significant analyst time to produce and was already out of date by the time it reached decision-makers. AI can make that analysis continuous, dynamic, and actionable.
- **Funding optimization and portfolio management** apply the same principle to the liability side of the balance sheet. Pattern recognition applied at scale across funding structures, pricing positions, and maturity profiles can meaningfully improve decision quality, particularly in volatile rate environments where static models lose relevance quickly.

The agentic shift is the ultimate goal

The furthest horizon is the move toward autonomous agents: AI systems capable of executing multi-step workflows without human intervention at each stage.

In automotive finance, this means operations where an agent can handle the full arc of a customer interaction, a contract event, or a fleet management decision, drawing on multiple data sources, making a series of conditional judgments, and completing the task without a human in the loop.

That's a meaningful shift from AI as a tool that assists decisions to AI as a system that makes and executes them. The product and customer value creation potential is real: new pricing models, dynamic contract structures, personalized leasing experiences that couldn't be operationally delivered at scale without autonomous capability. But the governance gap identified in chapter four applies directly here. Nearly three quarters of organizations plan to deploy agentic AI within two years but only one in five has a mature governance model for autonomous agents.

In automotive finance, where those agents would be operating on regulated credit decisions, financial transactions, and customer contracts, the regulatory dimension adds further weight. The EU AI Act explicitly classifies AI used in creditworthiness assessment as high-risk, with mandatory oversight and governance requirements that most organizations are not yet equipped to meet. The foundations, clean data, mature governance, stable infrastructure, and validated use cases in earlier tiers, need to come first.

The organizations investing in those foundations now are the ones who will be positioned to move when the governance frameworks catch up with the technology.

Recommendations: a roadmap for action

Recommendations: a roadmap for action

The research behind this report points in a consistent direction: the organizations making real progress on AI are not necessarily the ones with the largest budgets or the most ambitious roadmaps. They're the ones with the most structured approach. What follows is practical guidance for successful implementation of AI use cases in your organization.

Get the foundations right first

The most advanced players in our research didn't start with use cases but with infrastructure: a structured, accessible data layer, a governance model with clear ownership, and an architecture that wouldn't create lock-in as the tooling landscape continued to evolve.

That sequencing may feel slower at first, but it will provide competitive advantages going forward. Especially since data quality, governance framework, and technology stack flexibility are prerequisites for successful use cases.

For organizations still working through infrastructure migrations or operating on fragmented data environments, the priority is clear: finish that work before attempting to scale AI on top of it. Oversight frameworks need to be in place before autonomous systems go into production.

Besides, we've learned from our interviews that organizations who built governance early also move faster. Thanks to clear AI ownership structures, defined approval processes, and documented frameworks for human oversight, they don't wait for each use case to navigate the same approval cycle from scratch. They've done that work once, and they compound the benefit across every deployment that follows. Central steering like this also enables scaling and avoids duplicate efforts with compounding costs.

Separate automation from AI in your roadmap

Not every problem requires AI to fix it. One of the most practical things a leadership team can do is go through their AI roadmap and honestly categorize each use case: does this require non-deterministic inference, or does it require a well-designed rules engine and clean data?

The answer changes the investment case, the governance requirements, the timeline, and the risk profile. Automation problems solved with AI infrastructure add cost and complexity without adding capability. AI reserved for genuinely non-deterministic problems, where variability is high, data signals are complex, or decisions benefit from learned inference, delivers its full potential.

The frame from chapter five applies directly here: separate the two categories on your roadmap, resource them differently, and resist the pressure to call everything AI for the sake of strategic optics.

Know your quick wins before chasing complex deployments

High-effort, high-value use cases get the most attention, but they're also the ones most likely to stall. The effort-to-value ratio is a more useful organizing principle for a realistic roadmap.

There are genuine quick wins available in automotive finance: customer service automation, document processing, compliance monitoring, and basic operational routing. These are deployable on shorter timelines, carry lower governance risk, and generate real efficiency gains that build the internal credibility to fund the harder work. Identifying and sequencing those wins deliberately, rather than letting the roadmap be shaped by whatever generates the most excitement, is how organizations build production-grade systems rather than a graveyard of pilots.

Where to act yourself and where to rely on your vendor ecosystem?

The vendor question has changed. It is no longer primarily about features. It is about whether a platform exposes data reliably, supports the AI tooling an organization wants to bring, and can adapt as the landscape shifts. A platform operating across many customers also accumulates pattern recognition that no single organization can replicate, in contract behavior, vehicle operations, service costs, and residual value dynamics across markets. That aggregated intelligence is something no in-house build is positioned to match.

That reframing has a practical consequence for how organizations should think about their own AI investment.

Some use cases need to stay in-house. Retention modelling built on your customer history, commercial steering informed by your portfolio dynamics, and residual value strategy shaped by your specific asset mix are not candidates for vendor delegation. Done well, these capabilities reflect context that no platform solution can replicate, and they are where investment in internal AI talent pays its clearest dividend.

Building here is also about control. When an organization builds its own model on its own data, it can interrogate the outputs, understand the failure modes, and explain decisions to auditors and regulators. That transparency is not always available with packaged vendor features, and in a regulated environment it is not optional.

The answer looks different for use cases that are complex to build, non-differentiating in themselves, but operationally essential. Contract-aware automation, self-service execution, and the data exposure required for agent orchestration are not areas where most automotive finance organizations should be investing engineering capacity. They are the natural domain of the contract management platform.

When speed matters, the calculation shifts further. Not every capability warrants the lead time of an internal build, and organizations moving fastest are the ones recognizing that distinction rather than defaulting to one approach across the board.

The organizations that will be best positioned to move into genuine AI uses cases and, eventually, autonomous systems are the ones building on platforms that already have the right data underneath them. Choosing the right platform partners is, in that sense, one of the highest-leverage decisions available to automotive finance leaders right now.

Conclusion

Automotive finance has rarely been at the front of the pack when it comes to implementing new technology, and the same is true for AI. Nonetheless, a small number of organizations have already built serious foundations and are starting to generate results.

Others are still experimenting with scattered use cases, but it's not too late to regroup, restart from the data layer and switch gears. The use cases that will genuinely change the economics of automotive finance, from intelligent remarketing to commercial decisioning and fleet intelligence, remain largely untouched across the industry.

This means that organizations that invest in the right order now are the ones that will be positioned to act on it once the governance, the data infrastructure, and institutional confidence are in place.

For most of the industry's history, moving deliberately has been the right call. So far, it's been the best decision for AI, as well, but it's time to pick up pace now.

About SOFICO

SOFICO is the leading global provider of software solutions for automotive finance, leasing, fleet, and mobility management. Founded in 1988, the company serves customers in over 40 countries from offices in 10 locations worldwide. SOFICO's flagship product, SOFICO Miles Enterprise, supports the entire contract management lifecycle, helping organizations streamline operations, improve efficiency, and drive growth.

09 210 80 40

contact_global@sofico.global

www.sofico.global

SOFICO

SOFTWARE DRIVING MOBILITY