From GenAI to Agentic AI

Operating-model shift in financial services

WAVESTONE

Universität St.Gallen

Contents

1. Executive summary

Financial services is undergoing a structural shift driven by AI. Several pressures push banks and insurers toward AI at once. A shrinking, aging workforce makes automation necessary to sustain the work; boards push for fast, visible results under growing cost and margin pressure, in some cases mandating a start on agentic AI before the term has even been defined. The hype around generative and agentic AI, driven by competitors and vendors, raises the fear of being left behind. At the same time, regulation and legacy IT slow adoption. Within that tension sits agentic AI: a move from generative tools that improve isolated tasks to systems that orchestrate data, applications, and other agents into end-to-end workflows.

Based on thirty interviews across banks, insurers, and technology providers. This study leads to one conclusion: Operating-model maturity, not model capability, determines how far firms can scale agentic AI. Model capability already runs ahead of what most firms can absorb, so the gap that matters is organizational. The firms moving fastest are the ones that already had the process discipline and structure to put AI to work.

The key constraints are internal:

- DATA FRAGMENTATION
- MISSING BUSINESS–IT BRIDGING ROLES
- GOVERNANCE COMPLEXITY

A strategy derived top-down from business and IT goals, and backed at board level, sets the direction. The structure that follows is a hub and spoke model: a lean central hub, anchored in IT, that runs the core agentic work of composition, integration, and identity and access management for non-human actors, while the spokes deliver inside the business units. Centralize AI portfolio management in the hub to avoid duplication and enable reuse within a clear strategy. The model depends on people. Roles are shifting from doing the work to supervising and orchestrating it, so the transformation has to be carried with the workforce rather than imposed on it. Core domains such as claims, underwriting, and customer service will run as orchestrated value streams, with people moving from task execution to supervision. A composable enterprise ecosystem, internal and external modules combined freely, is a longer-term prospect, together with use cases no one can yet name. Reaching that stage depends on building the foundations now. The paper closes with recommendations for moving toward a more agentic organization.

2. Introduction and strategic context

Banking and insurance face high and conflicting pressure on AI adoption. The sector carries one of the heaviest regulatory loads of any industry and runs much of its core on decades-old legacy systems. A shrinking, aging workforce makes automation necessary; boards want fast, visible results under cost and margin pressure; and the hype around generative and agentic AI, driven by competitors and vendors, raises the fear of being left behind. Pulling the other way are two constraints that slow it down: a dense regulatory load, and an aging IT estate that limits what can be built.

The demographic strain is a sharp one in insurance. About one quarter of the German insurance workforce belongs to the baby-boomer generation, and a large share of them will retire within the next five years. The average employee age already sits just above 45 in both IT and back-office roles[1]. A shrinking workforce raises the pressure to automate processes, rethink service delivery, and hold operational capacity steady. It gives the AI agenda an urgency, and it changes how AI is seen inside the firm. Across the interviews, AI was rarely cast as a threat to jobs. It was framed as a way to keep core operations running as the workforce declines. This reframing matters in practice. It is the main argument executives use with works councils: AI as a way to absorb demographic pressure without cutting jobs.

Regulation is one of the main constraints. The sector works under one of the most demanding regulatory regimes of any industry. The Digital Operational Resilience Act (DORA), applicable since January 2025, sets ICT risk-management standards across twenty types of financial entity[2]. The EU AI Act classifies common financial-services applications as high-risk, including credit scoring and risk assessment and pricing in life and health insurance, and requires detailed documentation, human oversight, conformity assessments, and robust risk management[3]. These add to existing obligations under the GDPR, and national supervisory regimes. The combined burden slows the introduction of new technology, and the landscape is still fragmented and evolving. Firms must reconcile overlapping and sometimes contradictory mandates from European and national authorities, which makes coordination between institutions and regulators slow and complex.

[1]Gesamtverband der Deutschen Versicherungswirtschaft (2026). Demografischer Wandel: Versicherer setzen auf Digitalisierung GDV. https://www.gdv.de/gdv/themen/digitalisierung/demografischer-wandel-versicherer-setzen-auf-digitalisierung-197170

[2]European Parliament and Council (2022). Regulation (EU) 2022/2554 of the European Parliament and of the Council of 14 December 2022 on digital operational resilience for the financial sector. Official Journal of the European Union, L 333, 1–79. https://eur-lex.europa.eu/eli/reg/2022/2554/oj

[3]European Parliament (2025). Report on the impact of artificial intelligence on the financial sector (A10-0225/2025). https://www.europarl.europa.eu/doceo/document/A-10-2025-0225_EN.html

This pressure affects all competitors equally. No firm gains a lasting advantage by moving first into territory the regulator has not yet settled, and none falls behind by waiting. That is why the dominant posture is "stay prepared," not "move first." It sits in direct tension with the speed that AI now demands.

> **We process personal and highly sensitive data here. Technically, it's not the problem, but the regulatory effort, the AI Act, GDPR, DORA, that's where the biggest limitations are.**
>
> (Head of AI, German bank)

Both of these are shaped by a third, slower-moving force: the legacy IT estate that banks and insurers still run. These systems carry heavy technical debt. With most IT capacity tied up in maintenance, migration, and keeping systems running, the function has little room left for innovation, a pattern that matches the service-provider logic. AI ambitions expose this condition rather than create it, and it shapes where, and how fast, new applications can realistically be built.

A working definition of agentic AI as a systemic shift

Agentic AI systems consist of one or more AI agents that observe their environment, use available tools, and decide how to act toward a goal.

Drawing on advances in foundation models, they apply reasoning, perception, discovery, and planning to find ways to achieve their goals, using the memory and the tools available to them, namely data, applications, and other agentic systems[4].

The first-order agentic AI is an autonomous software entity whose six properties (goals, environment, memory, reasoning, actions, and tools) together enable goal-directed, adaptive behavior. Such an agent exhibits two dynamic capabilities, namely learning autonomy, updating its understanding from environmental input, and action autonomy, independently prioritizing and executing what it does[5].

Second-order agentic AI is the multi-agent case, functioning as a configuration in which coordination, delegation, and communication turn a set of autonomous agents into a coher-

[4]Li, M. M.. Reinhard, P., Leimeister, J. M., Catchword: Agentic AI. From Autonomous Agents to Enterprise Agentic AI Systems. Bus Inf Syst Eng, (2026).

[5]Li, M. M.. Reinhard, P., Leimeister, J. M., Catchword: Agentic AI. From Autonomous Agents to Enterprise Agentic AI Systems. Bus Inf Syst Eng, (2026).

[6]Li, M. M.. Reinhard, P., Leimeister, J. M., Catchword: Agentic AI. From Autonomous Agents to Enterprise Agentic AI Systems. Bus Inf Syst Eng, (2026).

ent system pursuing shared or distributed goals. At least one member must meet the first-order criteria, while the others may be non-agentic components or human actors who supervise, delegate, or collaborate within the configuration[6].

For financial services, the shift is beyond better text generation, moving toward orchestration that integrates departments, data, applications, and other agents into end-to-end work. It is this active execution that collides with the sector's defining constraints. Even when an agent is confined to a single team, agentic AI must operate within hard boundaries of rights authorization, governance, regulation, and policy the moment it can take actions, whether or not it ever leaves that team. Therefore, these boundaries rather than the technology itself determine the pace and extent of agentic AI adoption within a bank or an insurer.

Methodology

This study is based on a joint research initiative by Wavestone and the Institute of Information Systems and Digital Business at the University of St.Gallen (IWI-HSG). Wavestone, an international management and technology consultancy, provided access to senior decision-makers across financial services and kept the questions anchored in operational concerns. The University of St.Gallen, one of Europe's leading business universities with a long tradition of bridging research and practice, contributed the research method and an independent, evidence-led perspective.

The evidence base.

The study draws on 30 semi-structured interviews with banks, insurers, and selected technology providers. The interviewees come from four vantage points:

- Central AI units ("Hubs")
- Business-line teams ("Spokes")
- IT functions
- External providers

For most of the organizations we interviewed both the IT and the business side. That dual perspective is part of the design, much of what this paper reports is the gap between how the two sides read the same problem.

Who was interviewed.

Senior and technical roles from across the DACH financial sector:

- CIOs and Chief Data Officers; heads of AI, data, and innovation
- IT and enterprise-architecture leads; AI architects and data scientists
- Business and operations leaders

Organizations ranged from large cooperative and regional banks, retail-and-brokerage banks, to large composite and health insurers, smaller and recently merged carriers, and selected technology providers.

What was asked.

Standardized interview guideline ensured comparability. Core topics covered:

- AI organization and operating model (e.g., hub-and-spoke)
- Current maturity incl. ML, GenAI, agentic AI and production use cases
- Definition and understanding of “agentic AI”
- Roles, talent, and make-vs-buy decisions
- Governance, regulation, and co-determination
- Key blockers such as legacy systems, data, compute, sovereignty
- Outlook for the next ~2 years

Semi-structured format allowed flexibility to explore relevant topics in depth.

How the interviews were handled.

Every interview followed a common guide, was recorded and transcribed, and was then coded and analyzed systematically. Several researchers coded and analyzed the interviews across multiple review rounds.

Why the academic partnership matters.

The involvement of the Institute of Information Systems and Digital Business at the University of St.Gallen provides this study its rigor and its independence. As an academic institution, the university has no conflicting interest with any of the analyzed organizations.

Scope and limits.

This study is a qualitative assessment of the patterns, challenges, and developments shaping agentic AI in financial services. It captures a fast-moving field at a particular moment (interviews conducted between October 2025 and May 2026) and a DACH-centered sample.

3. Where financial services actually stand

3.1 The current state: mature ML, shallow GenAI, no agents in production

Classical machine learning is mature and widely deployed within our cases, especially in claims, fraud, and pricing. One large insurer in this study runs 75 such use cases in production. Its strongest example is a pipeline that screens about half a million legal judgments across millions of documents for recovery cases; it went live within six weeks and has recovered several million euros from fraudulent claims. The same insurer has pushed fraud detection beyond simple indicators into reading expert appraisal texts, and is rebuilding its fraud-network analysis so that an agent advances to the next step only once it has found enough red flags. A retail and brokerage bank describes its own AI in similar terms, traditional machine learning such as per-customer closing probabilities and next-best-offers, and places itself well short of agentic systems.

Generative AI is mostly limited to low-complexity use cases:

- **Summarization**
- **Drafting**
- **Data extraction**

Agentic AI is largely absent from production. The firm furthest along, a central AI Lab inside a large insurer, had not put a single agentic use case into production and described itself as "still waiting for the right one." The same leader with 75 live use cases rated his firm only three out of ten on readiness, the distance between doing AI and being changed by it.

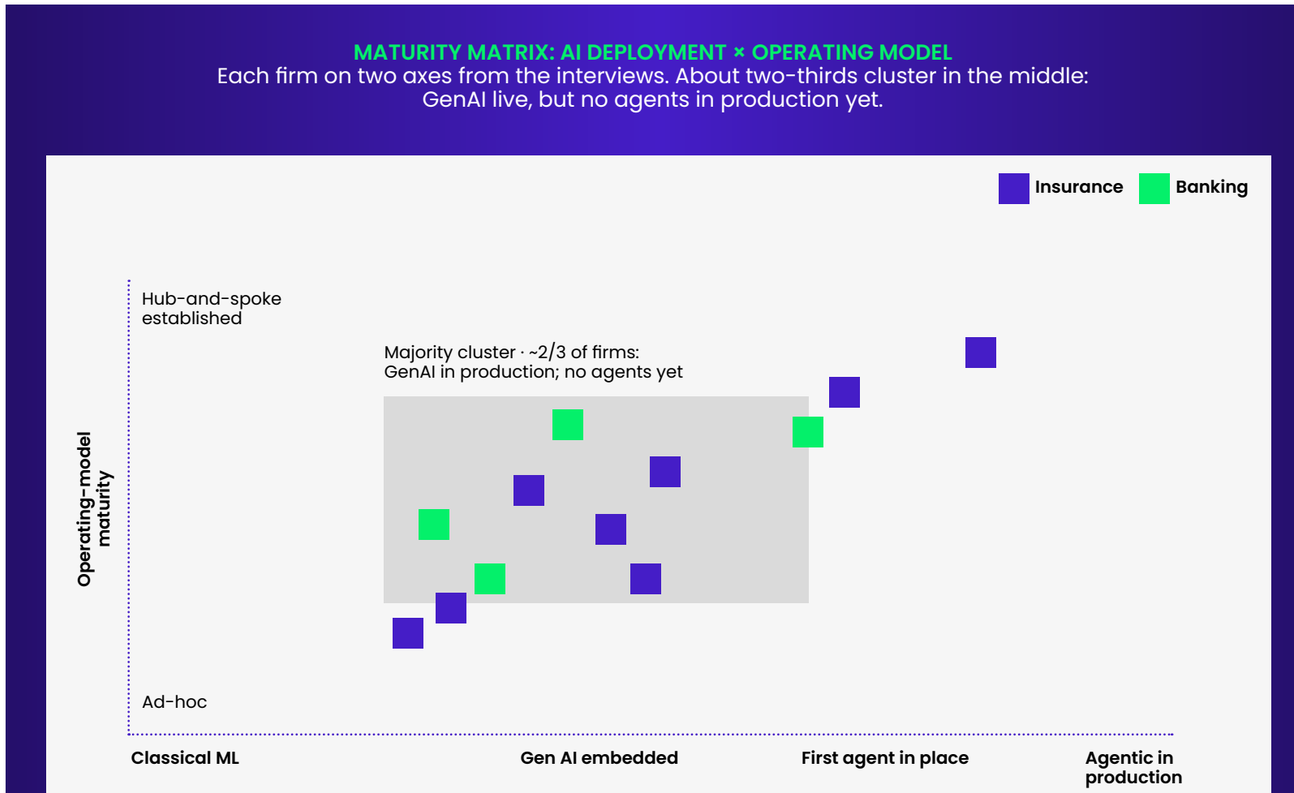

Fig. 1 Each interviewed firm placed on two axes. About two-thirds cluster in the middle, generative AI in production but no agents yet; a few laggards remain at classical ML, and a small leading group has first agents in place. Positions are indicative of the distribution the interviews describe, not exact scores.

An IT lead at a mid-sized insurer recalled the Big Data phase of a decade ago, much promoted and little delivered, and saw the same risk now for AI. AI lets a firm build isolated silos again, buying tools that each do something clever but only for themselves. Separate chatbots for sales, service, and claims, each fed its own data, are bound to give different answers because they do not draw on a common source. A provider made the same point, that AI does not work without clean data beneath it and that a fragmented data landscape resurfaces in the agents, which cannot pull from across sources. A small insurer's CIO put the foundation problem most plainly, admitting that "we don't even know our own data well enough," while wondering whether AI might help uncover the value hidden in it.

3.2 What separates firms: prior process scaffolding

The firms that move fastest are not the ones with the best technology. They are the ones that already had mature processes and structures. The clearest case is an insurer whose AI work sits on top of a 25-year-old Lean Six Sigma program with about 150 trained practitioners in the business. Its AI lead inherited that program and built on it, sending business problems through a structured improvement cycle in which AI is only one of several possible fixes, alongside process change, robotics, and ordinary software. That discipline avoids the common mistake of reaching for AI first and then searching for a problem to fit it. This lesson was learned the hard way. An early brainstorm produced 360 ideas, and about 350 had nothing to do with AI. The program did more than reorder priorities. It gave the firm a structured way to surface business problems, sort them, and judge which were worth pursuing, so candidates stayed manageable instead of piling up as an unfiltered wish list.

Starting from the business problem and running it through that structure is what filled the pipeline with work worth doing.

The scaffolding has different dimensions. One firm within our sample runs a weekly practitioner exchange and a reading group of PhD-level staff who track good agentic practice, the kind of capability a firm that has only bought a no-code platform cannot easily copy. The opposite problem appears at the smaller end of the market, where the challenge is not building the first solution but maintaining its quality. A small insurer's CIO had a university spin-off build a use case, then found that holding accuracy steady on the side was very hard. Each new model arrived twice as fast but with the hit rate down to 75 percent, and doing AI properly cost around 200,000 francs a year for just three modest use cases. What separates firms is the structure and discipline they built earlier, not how advanced their AI is today.

3.3 The use-case definition problem

The usual measure of progress, the number of use cases a firm has in production, rests on a unit that has no agreed definition. The first problem is that "use case" is applied to categorically different things. One firm's use case is an entire AI platform, another's is a single piece of AI software, a third's is a microservice operating at the level of one process step, and a fourth's is a single agent. Counts assembled from objects on such different scales cannot be compared, and the totals are inflated further by entries that are not AI at all. At one insurer an early list of 360 ideas held about 350 with no AI in them, and at another a catalog of 240 use cases proved to be mostly non-AI or long-standing work, much of it plain text recognition. The second problem is conceptual, because the term "agent" is itself undefined. "Agentic" is routinely attached to ordinary automation, and what a firm records as an agent use case is often a few chained prompts that summarize a report in a structured form, which is not an agent in any rigorous sense.

The effort behind a single use case varies just as widely, so the count fails even as a measure of volume: at one insurer a standalone prototype was finished in a day and a half, while threading the same idea through three systems that were being rebuilt at once took two years. The deeper issue is that counting activity reveals nothing about value. As one Chief Data Officer put it, counting use cases "says nothing about the value you could actually create if you rethought the business with what's now possible." This is why published adoption numbers cannot be taken at face value, and why the more useful question is where value actually is being realized.

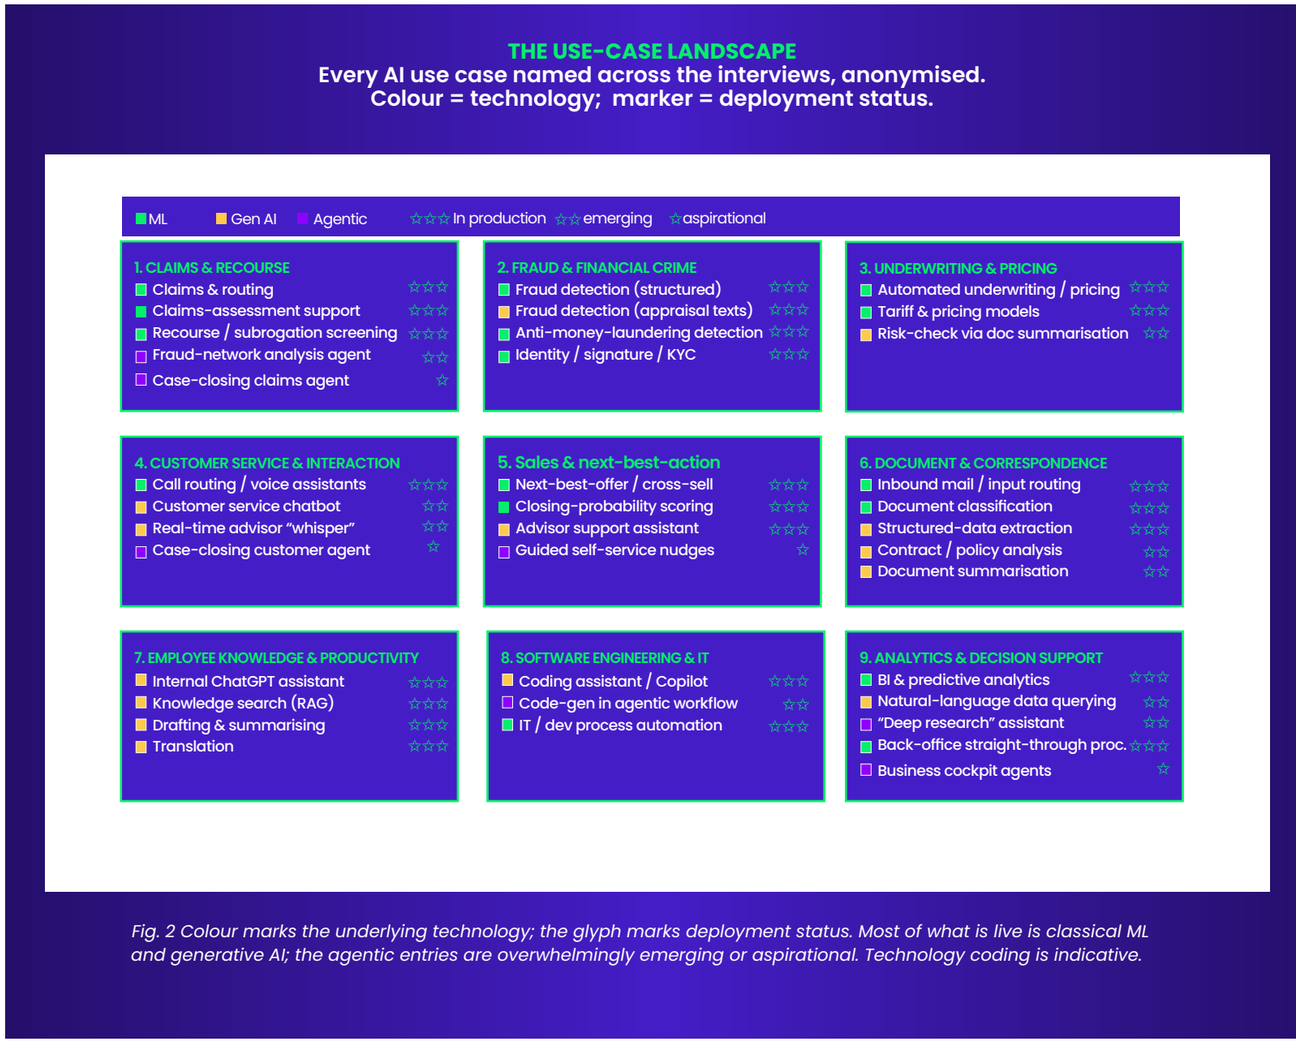

Fig. 2 Colour marks the underlying technology; the glyph marks deployment status. Most of what is live is classical ML and generative AI; the agentic entries are overwhelmingly emerging or aspirational. Technology coding is indicative.

3.4 Where value lands: the augmentation typology

The next question is where the value lands. Firms reach for customer-facing tools first, because the savings are visible and easy to claim. But the work that makes the front-end possible sits in the back office, and it gets too little attention. A head of operations at a large insurer said as much: his firm did not even have consistent data, and **"cleaning that up matters more than producing yet more use cases, because it's what makes all the others work better."**

There is a simple way to sort augmentation by how close it sits to the customer. At one end is direct customer-facing AI. Next is AI that supports the employee during the customer conversation, a three-way arrangement of system, employee, and customer in which the tool provides support to the person in real time. Then comes AI that improves back-end content which later feeds self-service. At the far end is back-end efficiency that reaches the customer only indirectly. Most of the real activity sits at the back-office end.

Customer-facing generative AI is rare in insurance because of a deliberate judgment about worse moments to hand to a machine. An IT lead at a public-sector insurer put it explains it with an anecdote:: a customer whose roof came off the day before, or whose parent has just gone into care, wants an answer, not a chatbot prefacing its reply with a note that the answers are not legally binding, and he doubted that outward-facing chatbots are the road an insurer wants to take. Banks draw a similar line when deciding where AI should be used. A head of AI at a retail and brokerage bank explained that the firm does execution-only, advice-free business, so the danger in any customer-facing agent is the moment it slips into giving investment advice, which keeps such plans at an early stage while the risk is still weighed. Even the cooperative bank that had built around fifty generative-AI use cases pointed its support assistant at advisors rather than customers, on the reasoning that advisors can judge the answers, and stayed cautious about letting the tools loose on customers. The ambition behind the caution is nonetheless clear.

3.5 IT and business see different worlds

By interviewing both IT and business leaders, in several cases within the same institution the study was able to examine that relationship directly. The central observation is that the two sides do not merely assign different priorities to a shared agenda; they operate with two distinct conceptions of what AI is and what it is for.

From the business side, AI is construed instrumentally: as a means to an end, defined by the problem it is deployed to solve rather than by the technology itself. The unit of concern is the outcome: a reduction in cost, a gain in revenue, a better-served customer and AI is one instrument among several for achieving it. On this view the governing question is never "what can the technology do?" but "which problem does it solve, and is that problem worth solving?" Business interviewees accordingly drew a sharp line between colleagues who treat AI as a tool for defined problems and those who pursue technical novelty for its own sake "using a cannon to shoot a sparrow" and located their own mandate firmly in the former.

From the IT side, AI is seen as foundational. Its value depends on the technical substrate beneath it: clean and accessible data, coherent interfaces, and an architecture in which systems can interoperate. Here the unit of concern is the system rather than the individual outcome, and the conception is marked by a demystified realism earned through earlier cycles of inflated expectation. As one IT lead observed after a decade in the field, the operative message carried into the business is simply that "AI is not magic." Where the business asks which problem AI solves, IT asks which conditions must hold before it can solve anything at all, treating data quality, integration, and reuse as the prior questions.

The clearest sign that these are two distinct conceptions of AI, the instrumental and the foundational, is that some firms have built them into their org charts. One bank deliberately split its AI unit into three groups. The first, "Business AI," is staffed with strategists, "the creative ones," who want the most advanced of everything. The second, "Data Science," exists to bring them back to reality: "but it won't work; let's see what will." The third, "AI Governance," connects both to compliance. The point is that the bank does not treat the gap between the ambitious and the realistic view as a problem to be fixed. It assigns people to each side and lets the disagreement do useful work.

The divergence, moreover, is not confined to priorities; it extends to definition. The two functions do not share a common understanding of what qualifies as "agentic": an initiative the business regards as conventional automation may be presented by the center as an agent, and the converse occurs with equal frequency. This unresolved question, what genuinely changes in the transition from generative to agentic AI is the subject of the following chapter.

4. From GenAI to agentic and what actually changes

The shift from generative to agentic AI is less about new capabilities than about composition. While GenAI improves single tasks such as summarizing, drafting, or coding, agentic AI connects models, data, and systems into end-to-end workflows. For banks and insurers, the real value therefore lies not in isolated use cases, but in the interfaces and interoperability that allow them to be combined.

In financial services the word "agent" is used for two different things, and the difference is rarely stated out loud, because everyone uses the same word. Underneath it sit two different ideas of what an agent is, what it is for, and what has to exist before it can work. As one insurer's AI lead put it, there is "no established definition yet, everyone understands something different," and yet the instruction from the top is to start with something agentic, with no clarity on what that means. Much of the confusion in the current debate, including inflated maturity figures, mislabeled use cases, and ordinary activity mistaken for transformation, comes from this split. Setting the two ideas side by side is the necessary first step, because the idea a firm starts from decides how its systems are built and how far they can scale.

Two competing framings of "agent"

The first framing is the process-bound agent: a narrow specialist embedded in one defined process step. It:

- **Takes a clear input and returns a clear output;**
- **Acts autonomously only within tight limits;**
- **Improves one task without changing the surrounding process;**
- **Is often used as an assistant, extraction component, or team-specific tool;**
- **Works best where data quality or governance maturity is still limited.**

In this form, the agent is a more capable version of conventional automation. Its value is concrete and immediate, but bounded by a single point of integration. Choosing this simpler setup is not a sign of low ambition. In many cases, it is the mature choice, especially where an autonomous, cross-system agent would add complexity without creating additional value.

The second framing is the process-liberated agent: an agent that is not tied to one process step, but can be recomposed across systems and departments. Its value comes from orchestration: several agents working together under a steering instance that coordinates them toward a broader goal.

This is more than a better assistant. It redraws the process around what has become possible. But it also requires foundations most banks and insurers do not yet have:

- **AI-ready data;**
- **Interoperability across historically siloed systems;**
- **Identity and access management for non-human actors;**
- **Governance models that can control autonomous action.**

For now, this model is still more aspiration than production reality. Where it appears, it is usually visible as intent: modular agents, pluggable architectures, or process redesign efforts that prepare the ground for future orchestration.

The two are best understood as the ends of a spectrum rather than a strict either/or, with a lone agent, then orchestrated clusters, then a fully systemic setup in between. Which label fits depends on who is looking. An operational unit may see conventional automation where the center promotes the same system as agentic, so the label says little about how the thing actually works. This is also why much of what financial services reports as "agentic" does not meet any strict definition.

Build modular, evaluate end-to-end

The two concepts are not competing approaches. They are two levels of the same operating logic: build modular, evaluate end to end.

Agentic systems should be built as small, specialized components rather than as one large monolith. A single large agent is hard to control, test and trust. A system of narrow, reliable modules is easier to build and scale. Functions such as parsing an address, validating an account number or checking a domain-specific field can each run as self-contained components that together carry the process end to end.

It is the practical path toward agentic capability. As one AI lead at a cooperative insurer explained, individual cases are built first and then combined until they automate larger parts of the process. The narrow agent is the building block for a larger system.

Value in an agentic system is created across the whole chain, so the unit of assessment has to be the end-to-end process, not the component. The discipline, then, is deliberately to separate the two planes: build small, but measure large. Collapsing them produces one of two familiar failures: an unmanageable monolith when everything is built as one, or a celebrated single-step optimization that leaves the process as slow as it ever was when everything is measured as parts.

Built and judged this way, an agentic system can advance well beyond where most firms stand today. But even done correctly, the move toward orchestration runs into a ceiling, and that ceiling is where the genuine blockers begin.

The limits of orchestration

The process-liberated vision points past today's modest agents, but it does not extend without limit, and the limit depends on what the system is for.

For production work, maximal connection is an anti-pattern. If every agent can call every other agent, the system risks loops, quality loss, rising cost and weak accountability. The target is different: contained, auditable clusters of agents, scoped to a use case and reused where relevant.

This boundary is not a technical limitation. It reflects the need to explain what the system did and why. In regulated industries, systems that cannot be audited and explained often cannot run at all. This is why firms focus on questions such as how to audit an orchestrator, which agent may call which other agent, and how data lineage becomes compliance-ready.

For work that is still exploring, the rule reverses. To find new use cases, a firm has to let agents connect freely and experiment, in a sandbox or a low-stakes zone where the cost of error is small. Here more connection is the point, not the danger. This is the familiar split in information systems between exploitation and exploration, the two sides of organizational ambidexterity: bounded and auditable where value is realized, open and unbounded where it is discovered. The interviews show both sides. Firms build fast, low-stakes prototypes and let teams experiment freely, while keeping the production side bounded and auditable. The two settings are not in conflict. They are designed for different use cases.

On the spectrum from simple automation to single agents, orchestrated clusters and fully systemic setups, banks and insurers still sit toward the lower end. The reason is structural, not technical. Model capability already exceeds what most firms can put into production. What holds agentic AI back is therefore rarely the model itself. It is the operating conditions around it.

5. The blockers

The main barriers to agentic AI are not model-related. They sit in three layers:

- **External: regulation**
- **Structural: compute & sovereignty**
- **Internal: data, people, operating model**

What the world imposes on a bank or insurer, what the market and available infrastructure place out of reach, and what the firm imposes on itself.
This order is deliberate. Each successive layer is increasingly shaped by the firm's own decisions and therefore becomes more decisive for its success. The outermost layer—and the one most often misunderstood—is regulation.

5.1 Regulation and codetermination

Regulation is usually described as a uniform drag, with every initiative taking longer and costing more in roughly equal measure. In reality it does not weigh on every process equally. It falls hardest on the core processes that drive much of the business, and barely at all on the peripheral ones. Because so much rides on them, they are the most tightly regulated and the hardest to change, even though that is exactly where automation would pay off most. So banks and insurers drift toward the small, low-risk cases that clear compliance easily but return little, while the highest-value work never makes it into production.

The difficulty has two sources. The first is the demand for explanation. Under frameworks such as DORA, the EU AI Act, and BCBS 239, a regulated bank or insurer must be able to show a supervisor the reasoning behind a decision and trace it back to the underlying data. For a fixed rule this is straightforward. For a model that cannot explain how it reached an answer it is much harder, and the bar climbs as the decision carries more weight. The second is that scrutiny does not scale down. A small improvement to a core process is held to the same standard of availability, audit, and review as the process it sits inside, however minor the change. Together these mean that the incremental approach is hardest to take on exactly the processes where it would help most, because those are the ones the framework guards most closely.

Codetermination produces the same pattern from a different direction. When a change affects employees, whether their roles, their workload, or their exposure to automated decisions, it has to be negotiated with the works council, and consent becomes harder to secure the more far-reaching the change. The council's stance varies widely across the industry, from outright resistance to active support, and what decides the direction is less the body itself than the way it is brought in. The underlying condition is demographic: where much of the workforce will retire within the next five years and teams are already short-staffed, AI reads as relief rather than as a threat, and resistance gives way to cooperation. The more useful way to see employee representation, along with data protection, legal, and information security, is not as an obstacle to be cleared but as a set of functions whose job is to weigh risk. Brought in early and as a matter of routine, they rarely become the blocker they are assumed to be.

5.2 Compute and sovereignty

Agentic AI in production needs computing power on a different scale than conventional automation did. Many banks and insurers conclude that this power must be hosted on-premise, within infrastructure they own and govern, because their compliance teams treat external providers as the higher risk. On-premise infrastructure built for conventional systems does not scale to an enterprise-wide agentic AI transformation. Closing that gap requires additional hardware that cannot be procured on European terms, since both the advanced chips and the legal control over them remain almost entirely outside Europe.

The sovereign response would close that gap by building the capacity in Europe and running the models on infrastructure the institution or the continent controls. A single US supplier, NVIDIA, accounts for most of the world's AI accelerators, reporting data center revenue of $193.7 billion in its 2026 fiscal year, and the same chips are crucial for Europe's AI data centers[7]. Because the United States can restrict the export of its most advanced semiconductors, a critical input to European AI infrastructure remains outside European control. Further, the US cloud providers AWS, Microsoft Azure, and Google Cloud together hold around 70 percent of the European market. Legal control does not follow the location of the data. The US CLOUD Act of 2018 lets US authorities compel a provider headquartered in the United States to hand over data even when it is stored in Europe, which puts that provider's obligations in direct conflict with European data-protection law[8]. In that regard,

[7]NVIDIA Corporation. (2026). Current report (Form 8-K). U.S. Securities and Exchange Commission. https://www.sec.gov/Archives/edgar/data/0001045810/000104581026000019/q4fy26pr.htm
[8]Gineikyte-Kanclere, V. et al., 2025, European Software and Cyber Dependencies, publication for the Committee on Industry, Research and Energy, Policy Department for Transformation, Innovation and Health, European Parliament, Luxembourg.

the sovereign-cloud offerings meant to resolve this fall short. European sovereignty today therefore describes, for the most part, US technology running on European ground.

The investment figures for future AI infrastructure also show differences between the US and EU markets. For example, the US venture Stargate alone has committed around half a trillion dollars to AI infrastructure over the next few years[9]. In contrast, the European Union's InvestAI program mobilizes just 200 billion euros[10], and even that sum depends on the same US chips. What must be seen as critical for European banks and insurers is the concentration of capacity, the speed of deployment, and US control of the supply chain to fully scale for an agentic AI transformation.

5.3 Inside the organization

LEGACY SYSTEMS

Old core systems constrain what can be built, but the real drag is continuous modernization programs that never finish and consume the engineering capacity AI needs. Building new components is manageable; integrating them into landscapes never designed for them is hard.

DATA

Agentic systems act on data directly, so it must be complete, correct, current, and available on demand across systems that historically siloed it. Weak data produces unreliable output regardless of the model. No single function owns data end to end, so the problem is accountability before it is technology.

PEOPLE

Bridging roles that span technology and business are scarce, and that scarcity drives the friction between IT and business. Roles are changing as the work shifts from model development toward engineering and integration.

ADOPTION

Demographic pressure from a shrinking, aging workforce makes automation necessary but does not overcome that resistance. Human-friendly automation designs the work around people. Reducing an employee to confirming machine output erodes motivation and quality, while combining human and machine deliberately keeps the role substantive and the result reliable.

GOVERNANCE

There is no clear strategy: firms hold AI roadmaps, but a roadmap is not a strategy, and what is missing is where AI fits, a time horizon, and conviction the approach will hold. There is no top-down sponsorship: board backing gives teams legitimacy to make control functions justify their objections, and without it legal, procurement, and information security slow every initiative.

THE ARC

The constraints range from external factors, such as regulation, compute availability, and digital sovereignty, to internal, self-imposed ones. Technology itself is not the limiting factor. The most fundamental constraint is the organisation's operating model: the one factor it fully controls and the one that ultimately determines whether external opportunities and capabilities can be translated into value.

[9]OpenAI. Announcing the Stargate Project. https://openai.com/index/announcing-the-stargate-project/, 2023.

[10]European Commission, EU Launches InvestAI Initiative to Mobilise €200 Billion of Investment in Artificial Intelligence, Press Release IP/25/467, 11 February 2025, https://ec.europa.eu/commission/presscorner/detail/en/ip_25_467.

6. The operating-model shift

6.1 The hub and spoke structure, and finding a path in the middle

Most banks and insurers in the interviews are converging on a central hub with spokes inside the business. The hub holds the shared capability, which is strategy and governance. The spokes handle delivery. They sit inside the business units, where the use cases and the domain knowledge are. In part this is rebranded innovation management, more technical and more decentralized than the classic innovation funnel. One insurer started its AI Lab at the end of 2020 with two people, a lead and a data scientist. They worked like a startup, where everyone held every role. They walked through the company offering themselves as an advisory team, began with low-hanging fruit, and built early extraction cases as reusable microservices, so a thing built once did not have to be built five times.

The model takes four recurring shapes, mainly depending on where governance, budget and works council involvement sit:

- **Hub-heavy:** the center owns governance, data and infrastructure, while spokes execute. This creates control and scale, but can reduce local ownership.

- **Spoke-heavy**: business units own much of the delivery and sometimes even codetermination. This increases ownership, but risks fragmentation.

- **Lab model:** a dedicated unit holds its own budget and steering committee. Depending on the firm, works council involvement may sit with the lab or remain with the spokes.

- **External setup:** innovation is placed outside the regulated core, for example in a digital subsidiary, venture arm or IT subsidiary.

These variants show that hub and spoke is not a fixed blueprint. It is a flexible arrangement that balances central capability with embedded delivery.

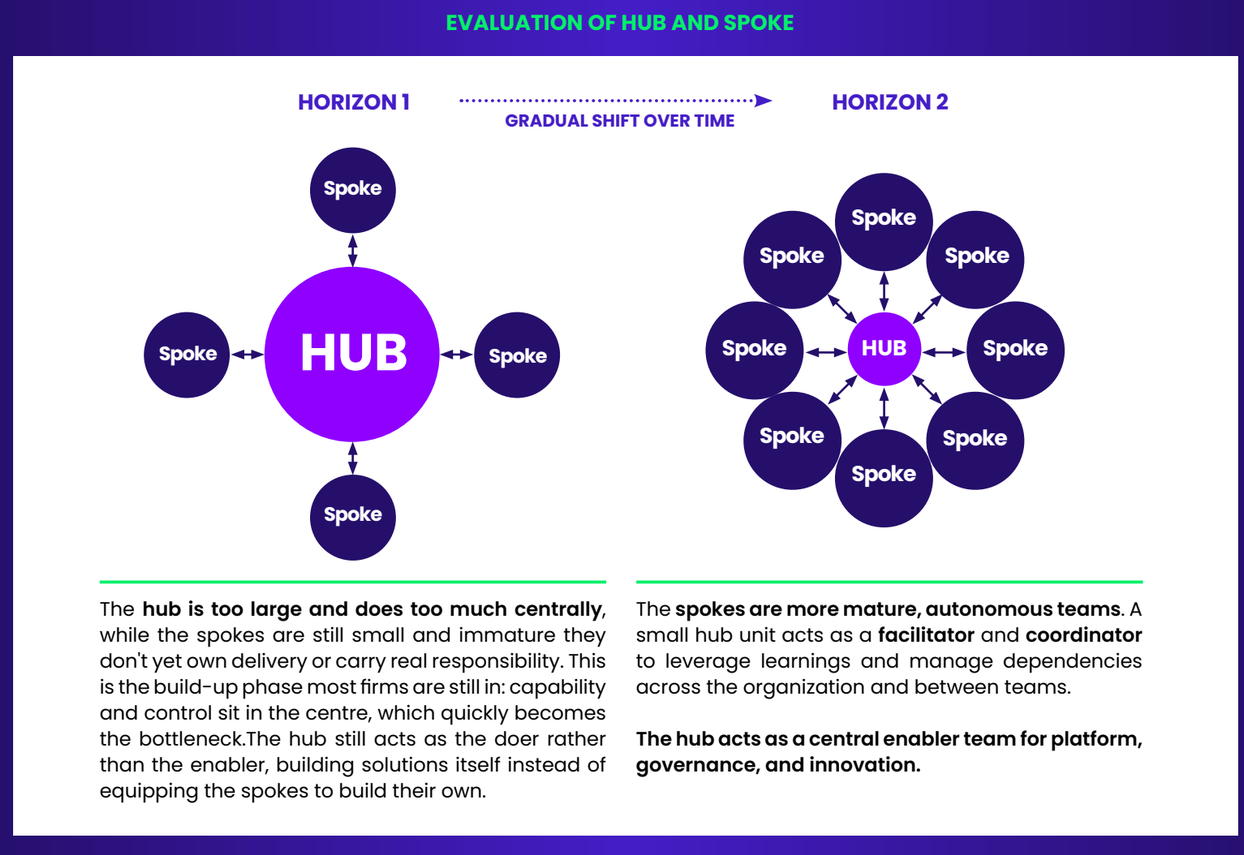

Fig. 3 The hub starts as an oversized doer and becomes a lean enabler as the spokes mature into autonomous, delivery-owning teams.

Two reasons explain why the model is still only partly built. The first is that the foundations are not yet in place. The data, the interfaces, and the ownership that agentic systems depend on are what most firms are still assembling. These are the blockers the earlier chapters set out. As the AI lead at one bank puts it, "in the end it is a transformation problem. The technology can already do more than we can put into production. Second, each pure form has a failure mode, which is why firms tend to converge on a hybrid middle over time rather than on full centralization or full decentralization. In the following we explain the main challenges.

Full centralization can break at the handover. A promising use case may stall once daily business takes priority, because the business has "other things to do than implement your AI use cases." Data ownership then becomes unclear. Structurally separate AI teams risk weak frontline support and may drift toward technical feasibility rather than business value.

Full decentralization can break the other way. Tools multiply with no shared goal and no guardrails. Units spend their own budget and buy separately. At one insurer, one unit wanted Cognigy, another wanted Parloa, and across three subsidiaries the firm bought all three in parallel.

The estate fragments until the same thing has been built several times. Spread this way, the work tends to cost technical depth, consistent governance, and robustness, and it overloads teams that carry daily operations and AI innovation at once.

The hybrid middle keeps a coordinating center without turning it into a bottleneck. That bottleneck appears when demand from the spokes outruns hub capacity, or when the hub sits too far from the operational context. The head of an AI Center of Excellence described his center as having stopped running every initiative itself. It now opens up and enables others while keeping only the coordinating role, "the bracket around it all." He knows what every team is doing without steering all of it. At his firm the split runs about 60 percent central and 40 percent decentral, with 30 to 40 units feeding one backlog ranked by value. Pushing everything into the spokes creates sprawl and forces a coordinating unit back into being, while building everything centrally would mean standing up a second IT, which is not realistic at scale.

The structure that most firms converge on, a coordinating hub with spokes embedded in the business, is largely agreed. What it does not settle are the organizational questions underneath it: where the central capability is located, how much should be centralized, and who owns governance and the budget. These questions cannot be deferred. They are central to the organizational design, because structure provides the framework, but governance determines who decides what. A large share of the work, including project direction, data strategy, change management, and IT integration, cannot be assigned cleanly to either the hub or the spokes, and unclear role boundaries are a leading source of failure in hybrid structures. The sections that follow take up these three allocation questions in turn.

6.2 Where the AI function should be anchored

There is no single right answer to where the AI function belongs. The firms in our sample have placed it in very different parts of the organization. The choice reflects the need to balance two objectives: running today's business efficiently while building new capabilities for the future. A bank or insurer has to run its regulated core efficiently, which is a task of exploitation, and at the same time build a capability that is still new and uncertain, which is a task of exploration. Where the function sits tilts it toward one or the other.

Most firms placed it in the business, because the early work was exploratory. The technical barrier was low, value was close to the customer, budgets often sat on the business side, and domain knowledge mattered more than engineering depth. The value was customer-facing, the budget sat on the business side, and the skill that mattered was domain knowledge, not engineering, since the value lay in knowing which problem to solve rather than in building the model. Several leaders told us the topic would have been "too far from the customer" if it had stayed purely in IT, above all where the IT organization was tied up in legacy work. Under those conditions the business was the natural owner.

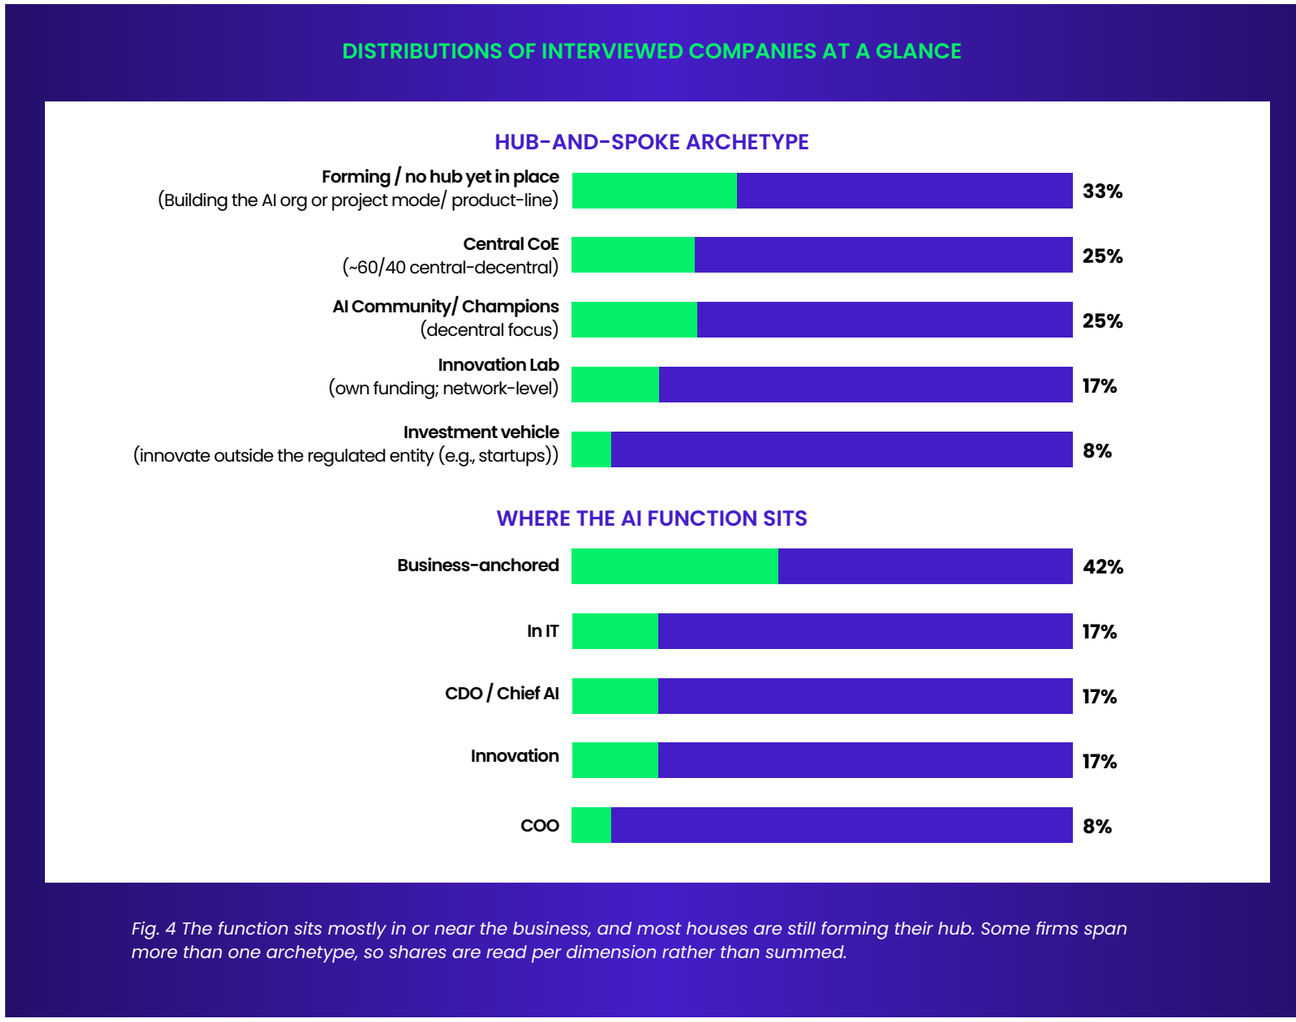

Fig. 4 The function sits mostly in or near the business, and most houses are still forming their hub. Some firms span more than one archetype, so shares are read per dimension rather than summed.

This placement creates friction when AI reaches IT only at handover. Three tensions recur:

- Risk is retrofitted: IT and compliance inherit undocumented data flows and controls after decisions are made.
- Silos multiply: use cases, tools and vendors spread across units before architecture is aligned.
- Lock-in increases: contracts with global cloud or model providers are signed before IT can assess long-term exposure.

The common cause is the entry point. Business teams commit to tools, vendors and spending before the functions accountable for architecture and risk are involved.

Agentic AI changes the nature of the work, and with it the logic of placement. With generative AI, the hard part sat on the business side: writing good prompts and getting people to adopt the tool. With agentic AI, the hard part moves. An agent does not just answer, it acts, so it has to be wired

into the systems where the work happens. The essential work becomes composition, orchestration, and integration across core systems, built on a clean data foundation and on identity and access management for non-human actors. Those core systems are the firm's policy administration, core banking, and payment platforms, which are old, heavily regulated, and security-critical. Connecting autonomous software into them safely is an integration problem, not a prompting or adoption problem. That is why integration has become the primary engineering challenge.

This is the argument for anchoring the AI function in IT. As the next crucial work becomes engineering, the function belongs where engineering, the platform, and the integration with core systems already sit. Here, the organisational anchoring becomes clear in the way the platform, the budget, and the engineering staff are housed within IT, with the function retaining decision-making authority over architecture, integration, and access.. By "IT" we mean the function that builds and runs platforms, not a back-office request desk. Anchoring in IT carries one risk. The function can drift into a downstream technical unit that only executes incoming requests and loses sight of where business value is created. The answer is not only to move it back into the business to stay close to value creation. Rather, it should be directed from the top, aligned with the firm's growth, cost, income, and risk objectives rather than driven by a queue of tickets.

A second view rejects the choice itself. Some leaders argue that the old IT department is the wrong home, because technology and infrastructure are a CIO topic while data and AI are mostly a business topic. They prefer interdisciplinary teams that hold business, IT, and the control skills of compliance, data protection, and security in one place. The two views agree on the substance. As agents move into core systems, the AI function needs two things at once: the engineering depth of a capable IT organization and direction set by the business. Anchoring in IT with business-set priorities is one way to get both. Interdisciplinary teams that combine the skills from the start are another. This is the same ambidexterity in a different shape: one capability that can both run the core and explore beyond it. One caution sits underneath both. Consolidating all procurement under IT to stop sprawl can also overload IT and slow adoption, so the consolidation has to come with capacity, not only with control.

6.3 AI strategy as the layer that steers the operating model

The operating model has many moving pieces: structure, sourcing, staffing, funding and governance. Without something above them to set the course, they drift apart. That is the role of an AI strategy, and its absence is itself one of the blockers; where no strategy sets the direction, the pieces, pull against each other.

A roadmap lists what will be built and when. A strategy sets a direction, and that direction gives teams a clear test for prioritization: a use case that does not serve one of the firm's priorities does not get built. The strategy starts with a realistic assessment of where the bank or insurer stands: its maturity, where AI can create value, how ready its data, skills, and organization are, and how much risk it is prepared to govern. Priorities then follow from the firm's business goals, namely growth, cost, income, and risk. One insurer admitted it had started building before doing this work, where management said "do cool things" and a request to "solve all my problems with AI" left the team with no basis to prioritize. Use cases gathered bottom-up produce activity without a clear aim. The same logic sets the measure of success. A strategy is judged by the value released when the business is rethought around what AI makes possible, not by the number of use cases in production. As the chief data officer at one insurer put it: "Implementing as many use cases as possible is not a KPI that says anything." Ultimately, success depends on whether the firm remains committed to the direction it has set. Several firms in our sample call their AI strategy something that must stay variable and adaptable, but a direction the firm already expects not to hold is not a strategy. Setting direction is top-down, and it has to meet the bottom-up reality of what is feasible and what the business needs.

The strategy also steers decisions that would otherwise be made unit by unit. The clearest is make-versus-buy — although in practice, the answer is usually make and buy:

- **Build** what differentiates the core business;
- **Buy** what is generic or likely to become standard platform functionality;
- **Use** external partners for flexible capacity.

This becomes strategic because of vendor lock-in. If a differentiating process depends on one provider's model or platform, the firm has outsourced part of its differentiation. No-code platforms can deepen that dependency, while pro-code and interoperable architectures preserve optionality. Left to individual units, make-versus-buy decisions create sprawl, duplication, and costly purchases with limited reuse.

The strategy must also define the people plan, which capabilities to build internally, hire, acquire or partner for. The roles are shifting. The pure data scientist is giving way to the data and AI engineer, as systems integration has become a core capability. The AI architect is the role most firms name as missing and most need, and at several it does not yet exist. New roles are emerging around the model, including AI champions in the business and the spoke leads who carry delivery.

A strategy also changes what governance can do. With an endorsed direction, the control functions weigh each request against a shared goal rather than refusing by default. That shift, from governance a brake to governance as an enabler, is the subject of the next section.

6.4 Governance: from brake to enabler

Governance in financial services is built to control risk, and without a clear direction it defaults to saying "no". Agentic systems make that default harder to escape. They are not deterministic in the way older software is, and they are not random either: an agent pursues a goal and decides how to reach it, so the same task can run differently from one time to the next, and the exact path is not fully predictable in advance. The hard question is therefore not how the output is produced but who is accountable when an agent acts on its own. Across the interviews, this open question of who owns the liability for an agent's action was named among the largest blockers. Governance built to approve a fixed, pre-specified behavior has less to grip onto when the behavior is decided at runtime.

A control function asked to clear an undefined risk against no stated priority. A clear direction removes that asymmetry. This does not move governance to the board or turn it into a strategy decision. It means the firm has agreed in advance what it is trying to achieve and how much risk it is willing to carry, so the control functions can judge each request against a stated objective and a stated tolerance instead of from nothing. Approval then turns on conditions rather than on permission. The sector describes this as moving from the red pen to the yellow one.

With that direction in place, governance becomes a shared decision across roles:

- **Business asks whether the use case is worth doing.**
- **Data science / IT checks whether it will run.**
- **Compliance and risk assess whether it is allowed.**

Each role effectively answers one question, is the use case worth doing, will it run, and is it allowed, which together cover desirability, feasibility, and admissibility. No single function decides alone. At one insurer, governance had, in the executive's words, "never refused a single case, only made more work of some." The result of a clear strategic direction is greater selectivity. Strategically aligned initiatives move forward more quickly, while low-value initiatives are rejected more decisively.

Two mechanisms put this into practice, and they sit at two levels: the firm's own internal rules, and the external legal requirements it has to meet. The internal mechanism bounds an agent's autonomy by the firm's own risk classification. One insurer runs a green, yellow, and red traffic-light system that scores every use case on two dimensions: how sensitive the data is, and where it is processed. The resulting cell says whether the use case may proceed and under what conditions. The external mechanism sets the legal floor and builds on top of it. It takes the EU AI Act as the baseline that every use case must clear, then wraps the firm's own risk-control framework around the deployed models and uses platform-level access control as the technical backstop.

Identity and access management was built for humans. Agents now need their own identities, their own permissions, and their own audit trails, so the firm can see what an agent did, when, and why. Several leaders called this the central question of agentic governance: what may an agent do, what may it not, and how is it monitored and traced. In multi-agent setups, it extends to which agent may call which other agent. It also addresses a security fear raised in the interviews, that an agent could be manipulated into passing information to people who should not see it. None of these mechanisms work without a clear direction above it, since each one presupposes an explicit view of which work carries priority and how much risk it justifies.

Put together, these mechanisms are what turn governance from a brake into an enabler. With a direction to weigh against, a risk classification that sorts use cases, a legal floor everyone clears, and identities and audit trails that make an agent's actions traceable, the control functions can say "yes" under stated conditions instead of defaulting to no. Low-risk work moves quickly, higher-risk work proceeds with the right controls, and only genuinely misaligned work is stopped. Governance still protects the firm, but it does so by enabling the right work rather than by blocking work in general. With that, the operating model is complete: a structure that organizes the work, a location that anchors it, a strategy that steers it, and a governance stance that enables it.

7. Scenarios for the future

Agentic AI is the composition of multiple agents into a system that carries work from end to end. The key question over the coming years is therefore one degree, namely how far that composition extends. For any given bank or insurer, the limit is set by the maturity of its organization and the extent to which it has addressed its key constraints, not by the quality of the models, which already exceed what most firms can effectively absorb.

7.1 The timeline

Last year was about organizational design, institutionalizing Hubs, Labs, and Centers of Excellence. This year, where most firms sit, is about compliance and the organization catching up, building foundations, negotiating works council approvals that any change affecting employees requires, and selecting high-ROI use cases worth doing. Practitioners described themselves as "in the starting blocks, ready before the party starts, still checking the rules of the game," and several named 2027 as the point by which they expect to be working with agentic systems in earnest. Next year and beyond is about making the work systemic, with orchestration as the source of value and multi-component systems replacing one-off tools.

A second sign of maturing is that the more deliberate firms have started to plan against scenarios rather than react to the market and to technical developments. One insurer's AI lead intends to build his strategy on a published foresight study that maps how AI bears on technology, money, people, skills, and processes, in place of the common habit of "jumping a meter forward, like everyone else." The three scenarios that follow are a version of that exercise: three depths of composition that agentic AI in financial services can reach, and what each asks of the firm.

7.2 Agentic AI scenarios for the next years

From the interviews we derived three scenarios over coming years. They describe where the firms in the sample are heading and what the future could look like. The human role does not disappear as agents grow more capable. It moves up. In the first scenario, the human supervises, in the second the human steers, in the third the human orchestrates, holding the relationships and the decisions the agents cannot. A provider put the underlying reason simply: value is created where agents cannot stand on their own, in human interaction and personal relationships. Deliberately defining the role of humans within that process makes higher levels of agent autonomy safe to adopt.

Scenario A: Augmentation.
This is first-order agentic AI: a single agent on a single task, supervised. The assistants behave much like robotic process automation, summarizing, drafting, routing, and answering from a retrieval knowledge base. They deliver productivity gains without structural change, and the person stays in the loop as supervisor. This is where most firms operate today, and for risk-averse functions it is a stable endpoint rather than a way station; a cooperative bank's innovation lead likened where these agents are heading to "RPA 2.0 or 3.0." The demand on the organization is minimal, which is why augmentation is reachable now.

Scenario B: Domain composition.
This is second-order agentic AI: several agents coordinated end to end inside a single value stream, such as claims intake, recourse, or underwriting support, run back-office and mid-office first and customer-facing last. It is where most of the sample is heading. The first composed workflows appear within about a year, but turning a whole value stream over to agents takes longer, on a three-to-five-year curve that is slow at first and steep at the end, because the customer-facing steps come last. A provider described the shift as people ceasing to do the work and beginning to steer it through multi-agent systems. Customer-facing comes last for a reason of risk: the storm-damaged-roof moment in insurance, and the line between guidance and licensed advice in banking. The human role moves from supervising single tasks to steering the whole stream: setting its goals, watching it run, and handling the cases the agents escalate.

Scenario C: The composable ecosystem.

Agents and whole multi-agent systems composed freely and recomposed for each use case, reaching beyond the firm in two directions. The first is across organizations, where the firm's agents work with those of its partners and technology providers. The second is out to the open internet, where agents discover and coordinate with agents the firm does not control. It demands the complete operating model: an enterprise platform, interoperability, identity and access for non-human actors, and the scarce AI-architect role. Its promise is the use cases no one can yet name, the ones that become possible only once the pieces combine freely. The interviews are consistent that this state is gated not by model quality but by foundations no firm in the sample yet holds: AI-ready data, interoperability, the missing AI-architect role, identity and access for non-human actors, and compute. A Chief Data Officer listed exactly those gaps and concluded that the limit is organizational maturity, "not a technical limitation."

The interoperability piece deserves its own note, because it is the part moving fastest. A composable ecosystem needs agents built by different vendors and frameworks to discover, understand, and work with one another, and the standards for that are being built in the open. Alongside the cross-vendor semantic layer that providers described, two protocols are emerging as the wiring of multi-agent systems: Anthropic's Model Context Protocol, which connects an agent to tools and data, and Google's Agent2Agent protocol, which lets agents discover and coordinate with each other across platforms. Agent2Agent was placed under neutral governance at the Linux Foundation in 2025, backed by the major cloud and enterprise vendors and positioned as complementary to the Model Context Protocol. The horizontal agent-to-agent interaction that the composable ecosystem assumes is therefore no longer hypothetical; it is being standardized now, which is one reason a two-to-three-year horizon for the first cross-system ecosystems is plausible. For banks and insurers, the scenario a firm reaches depends on how deliberately it builds the foundations now, and on what it decides agentic AI is worth. The reward at the upper end is a financial services business recomposed around agents, with claims, underwriting, and servicing run as orchestrated value streams, opening use cases that cannot yet be named; the banks and insurers that begin laying that groundwork today are the ones that will define them.

8. Recommendations

The five recommendations below build on each other in sequence. A strategy sets the direction for the organization. The organization then takes shape around it as a hub and spoke model, establishing a lean hub with strong responsibilities in the spokes. The hub runs the portfolio that keeps the work coherent. Governance turns that work from blocked to enabled. People carry the transformation through.

1

Start with an AI strategy that makes the vision explicit and wins board backing. Set the direction before the structure. Derive priorities top-down from the goals the firm already holds, which are growth, cost, income, and risk. For each goal, name where agentic AI creates value, so the firm has one vision to build toward rather than scattered initiatives. Secure explicit sponsorship for that strategy at board level. It gives teams the standing to act, unlocks budget, and brings the patience to push changes through. Give every team a clear basis to decline work that does not fit. Measure progress by the value released when a process is redesigned around what agents make possible, not by the number of use cases in production. The strategy is the top-down anchor that everything else hangs from.

2

Let the hub and spoke model follow from the strategy, and anchor the hub in IT. Translate the strategy into an operating model. A central hub owns the platform, the engineering and the governance. The spokes deliver inside the business. Anchor the hub in IT, where the binding agentic work now lives, which is composition, integration, and identity and access for non-human actors. Keep the hub lean, and set its direction from the top rather than from a queue of requests. Build the spoke network out deliberately over the coming years rather than all at once. Give each spoke the hub's targets, so delivery and ownership spread across the firm and daily work does not crowd AI out. Treat this as the firm's durable AI organization, not a temporary project structure.

3

Run AI portfolio management from the hub within IT. Stand up one portfolio inside the hub, ranked by value, and steer all AI work through it. The hub then becomes the central place where solutions and best practices are collected and reused. Route every initiative through a single backlog, and give the hub visibility over what each unit builds and buys, so the same use case is not built several times across the firm. Establish one definition of a use case and apply it consistently, so progress can be compared and priorities set on a common basis. Rank the backlog by business value and risk in line with the strategy. Make it the point where build versus buy is decided and where scarce capacity and budget are allocated.

4

Make governance an enabler by bringing the stakeholders together. With the strategy in place, set governance up to enable. Give the control functions an endorsed direction to weigh each request against, so a clear yes becomes possible. Bring the relevant functions into one decision early, before the prototype. These are the business, data science, risk, compliance, and the works council. The liability and codetermination questions take longer to resolve than the technical build, so they have to start sooner. Tie an agent's autonomy to the firm's own risk classes, and set the EU AI Act as the floor to build on. One insurer runs a green, yellow, and red matrix over data classification and processing location, maps each use case into a cell with the conditions for clearing it, and uses the matrix to steer the talks with the control functions, which it describes as fighting house to house. Run governance to prevent sprawl and to clear good work, not simply to slow it.

5

Drive the transformation through people, because nothing else carries it. Agentic AI changes what people do. Roles shift from performing the work to supervising, steering, and orchestrating it, and the skills that matter shift accordingly. Build and staff the bridging roles that span business and IT, above all the AI architects and translators on whom scaling depends. This capability matters more than adding further data scientists. Build AI literacy across the business, so the people who oversee agents can challenge and own their output. Bring the workforce along deliberately. Frame AI as relief for a shrinking and aging workforce, design the automation around the people who run it, and engage the works council early. The transformation then happens with the organization rather than against it.

Authors:

Eugen Ensinger
Associate Partner, Hamburg
AI as a Service
Next Generation Financial Services (NFS)

Luc-Phillip Heidel
Senior Consultant, Hamburg
Agentic AI Strategy & Transformation

Prof. Dr. Jan Marco Leimeister
Managing Director of the Institute of Information Systems and Digital Business at the University of St. Gallen, Switzerland and Chair for Information Systems at the Research Center for Information Systems Design at the University of Kassel, Germany.

Prof. Dr. Mahei Manhai Li
Assistant Professor for Management of Generative und Agentic KI in Organizations at the University of St. Gallen, Switzerland.

Teresa Grauer, M.Sc.
Associated research assistant at the Institute of Information Systems and Digital Business at the University of St. Gallen, Switzerland and research assistant at Research Center for Information Systems Design at the University of Kassel, Germany.

WAVESTONE

ABOUT WAVESTONE

Born from the rise of digital innovation, Wavestone has evolved into a leading consulting firm positioned at the convergence of business and technology. It helps the world's leading organizations scale new ways of working, transform operations and reinvent business models in the age of AI.

With 6,000 employees across Europe and North America, Wavestone combines business, transformation and technology capabilities to deliver end-to-end consulting services. AI has been embedded across the firm's expertise and ways of delivering services, enabling clients to translate its potential into tangible and sustainable value.

Wavestone is listed on Euronext Paris and has been certified as a Great Place to Work®.