SUTHERLAND®

WHITEPAPER

Agentic AI in Financial Services: Smarter Decisions, Real Outcomes

Executive Summary

Agentic AI reflects the next operating model shift for banking. It moves automation from scripted tasks and GenAI copilots into goal-driven decisioning that can decide and act across end-to-end processes like onboarding, lending, disputes, payments servicing and financial crime operations. This shift matters because the industry is running into hard constraints: persistent cost pressure, exception-heavy work that does not scale with traditional automation and rising governance expectations around transparency, resilience and third-party risk.

Two market signals define the moment. Gartner expects **40% of enterprise applications** to include **task-specific AI agents by the end of 2026**, up from less than 5% in 2025.[1] Gartner also predicts **over 40% of agentic AI projects will be cancelled by the end of 2027** because of unclear value, poor governance and weak data readiness.[2]

Taken together these signals say something practical: agentic AI is moving quickly into the stack but many institutions will waste money unless they treat agents like a controlled digital workforce with auditable evidence and measurable outcomes.

This paper introduces an original framework to scale agentic AI in BFS without creating model risk debt: the **TRUST Loop**. It also proposes a portfolio lens to select the first domains to industrialize, and it defines what a BFS-grade AI Hub must deliver to move from pilots to production: orchestration, embedded guardrails, evidence-by-default and reusable agent patterns.

1 Gartner, “Gartner Predicts 40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026, Up From Less Than 5% in 2025” (press release, Aug 26 2025).
2 Gartner, “Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027” (press release, Jun 25 2025).

1. Why agentic AI now

Banking has always been a scale business. Scale used to come from branch footprints, low-cost deposits and straight-through processing. Over the last five years the physics changed.

- **Revenue headwinds increase the premium on productivity.** A more modest banking environment in 2026 heightens focus on operational discipline and cost-to-serve.
- **Customer and regulator expectations keep rising.** Faster resolution, fewer errors, consistent explanations and resilience are no longer differentiators.
- **The work is still exception-heavy.** Many core journeys still require humans to navigate multiple systems, interpret documents, reconcile policy and document decisions.

GenAI copilots help with content and knowledge retrieval but copilots mostly stop at advice. Agentic AI is the next step because it can execute within bounded autonomy: interpret intent, retrieve context, decide under policy, act in systems and generate an auditable case file.

McKinsey's banking research describes the direction clearly: multiagent systems can become "virtual coworkers" that plan and execute tasks and workflows like credit underwriting when combined with enablers like retrieval, intent detection, summarization and safety guardrails.[3]

3 McKinsey, "Extracting value from AI in banking: Rewiring the enterprise" (2024).

2. What “agentic” means for financial institutions

An AI agent is a purpose-built digital worker that can interpret intent, retrieve context, decide within policy, act through tools and systems and document evidence and rationale. Agentic AI is the orchestration layer that coordinates multiple agents, enterprise tools, data sources and human controls to complete end-to-end work.

Here’s how automation evolved over the years

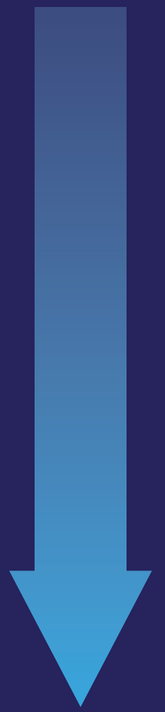

Evolution Wave	Primary capability	Typical use	Where it breaks
Rules and scripts	Deterministic steps	Data entry, routing, reconciliations	Fails in exceptions and policy change
RPA at scale	UI task automation	Swivel-chair reduction	Brittle integrations and limited reasoning
GenAI assistant	Language generation and summarization	Knowledge search, notes, drafting	Stops at recommendations, cannot reliably execute actions
Agentic AI	Goal-driven decision + action with evidence	Disputes, onboarding, KYC refresh, fraud triage, collections	Requires orchestration, governance and data readiness

Table 1. Evolution of Automation in Financial Services

3. Market signals: Fast adoption and Predictable failure modes

Three signals matter for leaders across financial institutions:

SIGNAL 1

Agents are becoming embedded features. Gartner's forecast that 40% of enterprise apps will include task-specific agents by end of 2026 means agentic capability will become a default platform feature not a bespoke innovation program.[1]

SIGNAL 2

Adoption is accelerating. PwC's May 2025 AI Agent Survey reports 79% of executives say AI agents are already being adopted in their companies and 88% plan to increase AI budgets in the next 12 months due to agentic AI.[4]

SIGNAL 3

Many programs will not survive production reality. Gartner predicts over 40% of agentic AI projects will be canceled by the end of 2027.[2] For BFS, that failure rate is not surprising. Agents touch regulated decisions, customer outcomes and money movement. Weak controls turn into reputational risk faster than most industries.

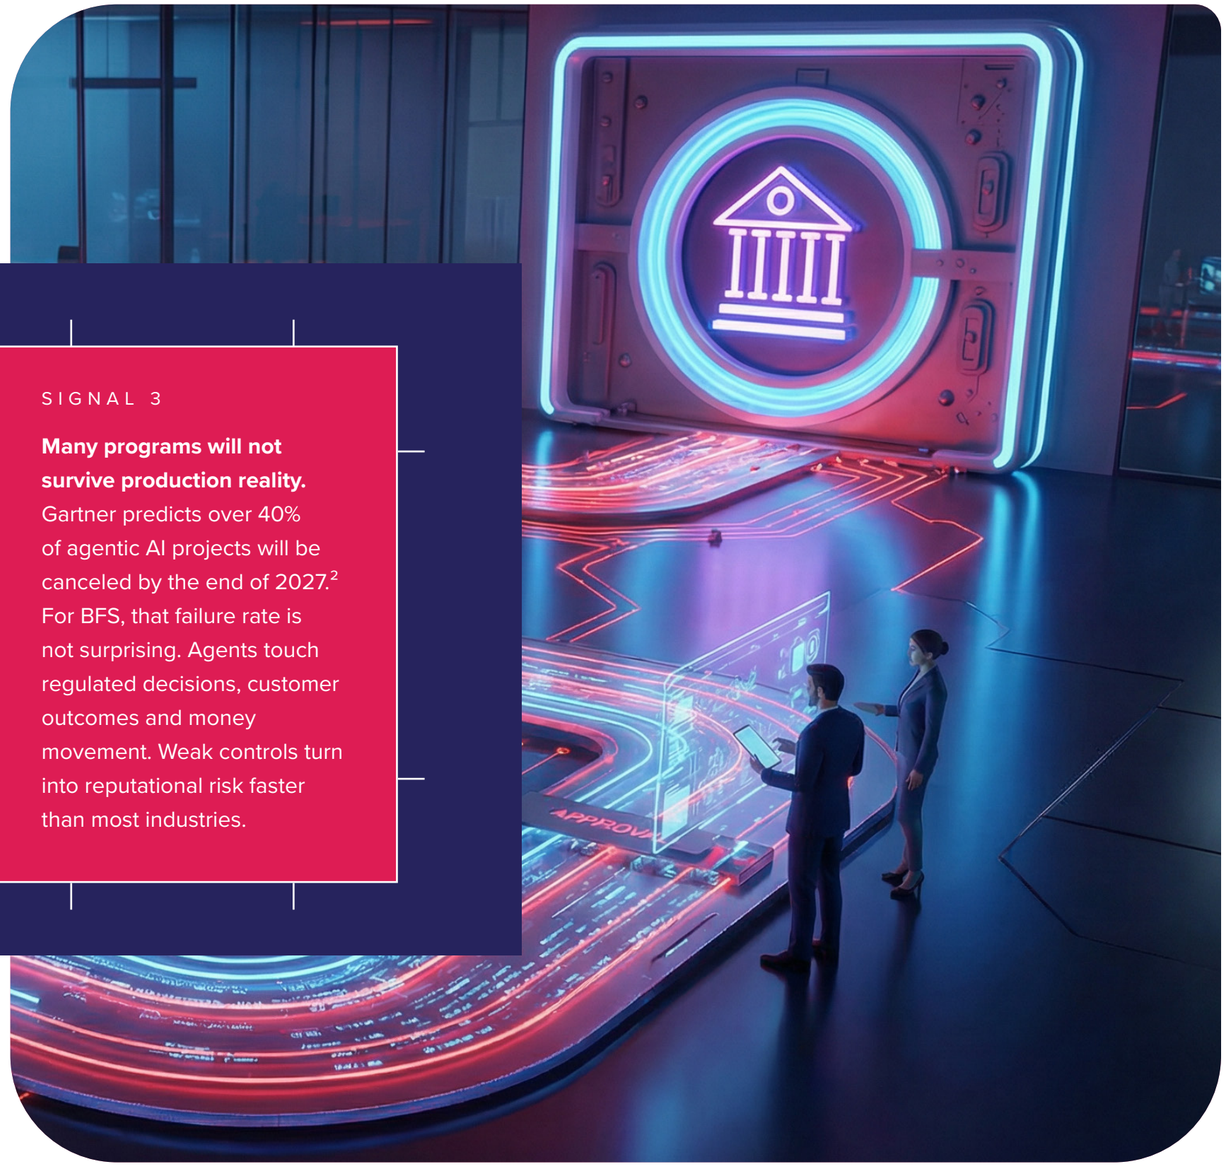

Implication: The value gap will widen between banks that industrialize agentic AI with governance and banks that run pilots without a production operating model.

1 Gartner, "Gartner Predicts 40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026, Up From Less Than 5% in 2025" (press release, Aug 26 2025).
2 Gartner, "Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027" (press release, Jun 25 2025).
4 PwC, "AI agent survey" (May 2025).

4. Trust is the constraint and the differentiator

Agentic AI expands the number of “micro-decisions” made inside operations: who gets onboarded, which disputes are accepted, what gets escalated in financial crime operations and what remediation actions occur in servicing.

Supervisors and regulators will ask three questions:

1. **What did the agent do**
2. **Why did it do it**
3. **Can you prove it used the right evidence and policy**

Recent regulatory posture reinforces the direction. The Bank of England and FCA’s 2024 survey reflects broad AI usage in financial services and a supervisory focus on governance as adoption scales.[5] The EU’s Digital Operational Resilience Act sets expectations around ICT risk management and oversight of critical third parties.[6] Even for US banks that do not operate under DORA, the discipline is instructive: resilience and oversight must be engineered into the operating model, not bolted on.

Agentic AI increases trust requirements because it is closer to execution. That is why auditability, traceability, access control and human-in-the-loop are not optional. They are the adoption unlock.

5 Bank of England and Financial Conduct Authority, “Artificial intelligence in UK financial services 2024” (2024).
6 European Banking Authority, “Digital Operational Resilience Act (DORA)” overview.

5. The TRUST Loop

To scale agentic AI responsibly, institutions need a closed-loop governance model.

Sutherland's TRUST Loop ensures that autonomy is matched by accountability at every layer of execution. It creates a feedback system where performance data, risk signals and audit findings continuously refine agent behavior. That feedback loop is what transforms agentic AI from a productivity tool into resilient infrastructure.

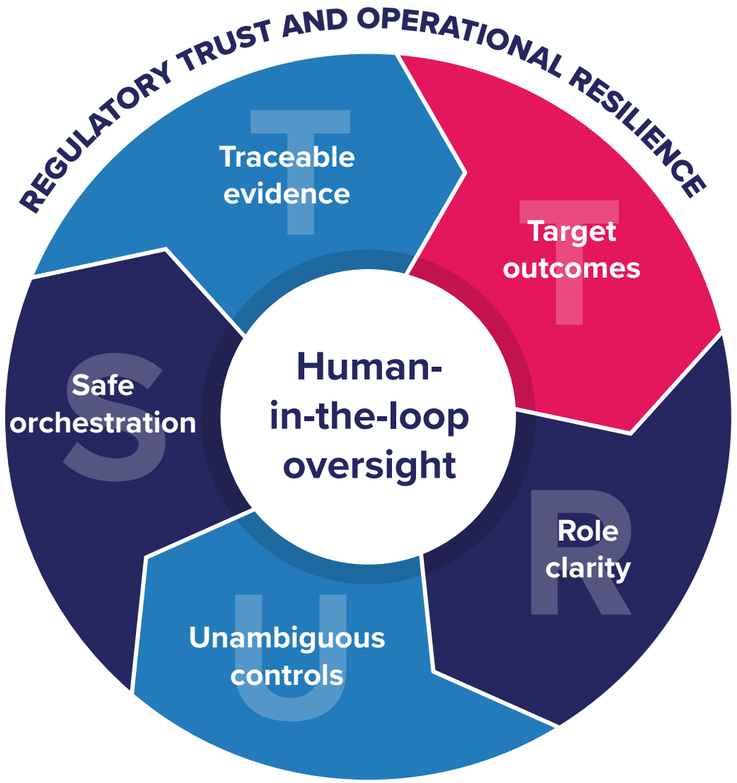

Sutherland TRUST Loop for Agentic AI

T: Target outcomes
Define value in business units: minutes saved per case, cycle time reduction, fewer reopen rates, lower complaint volume, fewer compliance exceptions. Assign each agent one primary outcome and one primary risk metric.

R: Role clarity
Treat agents like job roles. Define what the agent owns, what it recommends and what it executes. Publish boundaries for operations, risk and audit.

U: Unambiguous controls
Encode policy, thresholds and stopping rules. Controls include data access constraints, action permissions and escalation triggers for low confidence or novel scenarios.

S: Safe orchestration
Use orchestration to enforce sequence, approvals, separation of duties and exception handling. Orchestration is where you prevent runaway automation and ensure humans are pulled in at the right moment.

T: Traceable evidence
Log every prompt, context source, tool call, decision and action. Produce an auditable case file that mirrors human investigator documentation.

Most AI programs fail because they treat governance as review-based. The TRUST Loop makes governance architectural. Each element reinforces the others.

The TRUST Loop is not a checklist. It is an operating discipline.

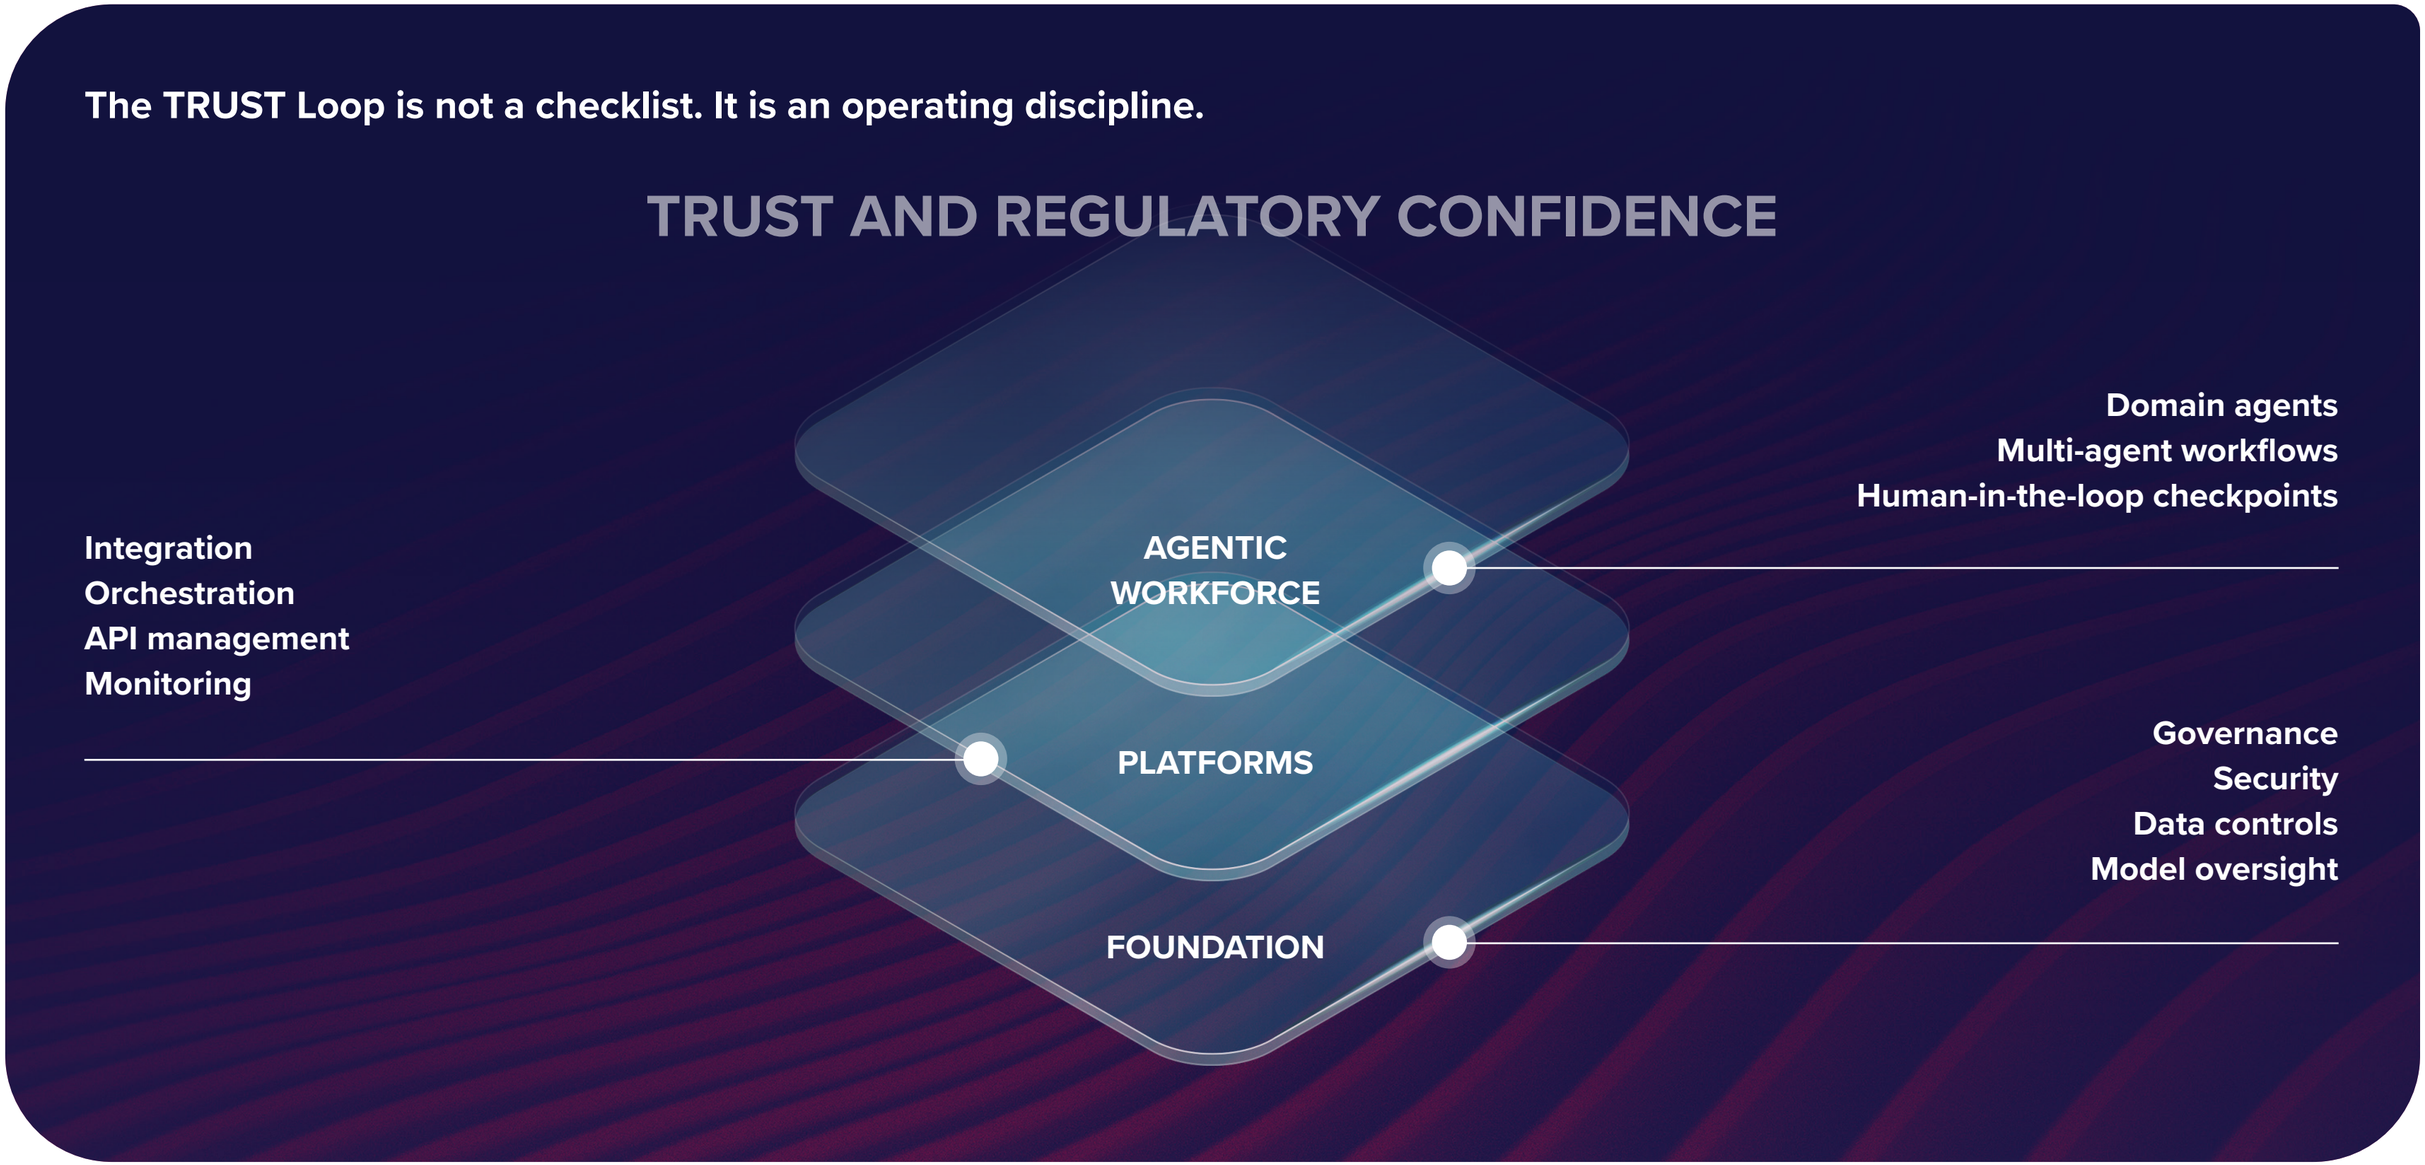

6. Where agentic AI creates outsized value in banking

The fastest path to value is not one giant agent. It is a portfolio of narrow agents that can be composed into multiagent workflows.

A practical selection filter is a two-axis test:

- **Operational friction:** high volumes, heavy manual navigation, repeated rework
- **Decision complexity:** repeated decisions with bounded policy and evidence

Domains that score high on both axes offer the strongest early return for agentic deployment. They provide measurable economics, controlled autonomy and defensible governance.

Agentic AI creates the most value when deployed in domains where operational friction and policy-bound decisioning intersect. The table below identifies high-impact banking domains that meet those criteria.

McKinsey highlights that value comes from rewiring domains rather than deploying isolated use cases and multiagent systems can support complex workflows when paired with enablers and safety elements.[3]

Domain	Why agentic fits	Value metric	Risk metric
Onboarding and KYC	Document heavy, policy-driven, external verification	Cycle time, manual touches per case	KYC exceptions, audit findings
Disputes and chargebacks	High volume queues, strict timelines, evidence handling	Cases per FTE, response time, complaint rate	Deadline breaches, incorrect outcomes
Financial crime operations	Alert volumes, repeated evidence review	Alerts auto-disposed, investigator time per true case	False negatives, drift
Collections and hardship	Segmentation and treatment paths can be personalized with controls	Cure rate, cost per dollar collected	Fair treatment and conduct risk
Servicing	AHT driven by system motion and knowledge retrieval	AHT reduction, containment, first contact resolution	Incorrect advice, privacy incidents

Table 2. High value agentic domains and what to measure

3 McKinsey, “Extracting value from AI in banking: Rewiring the enterprise” (2024).

7. What BFS-grade agents look like in practice

A common misconception: “agentic” equals smarter chatbot. In production banking, a BFS-grade agent follows a repeatable pattern:

Ingest → Validate → Decide → Act → Document

This sequence mirrors how human analysts operate. The difference is that the agent executes within defined policy boundaries, system permissions and governance controls.

What distinguishes agentic systems from traditional automation is composability. Rather than a single monolithic model attempting to handle an entire workflow, production-grade architectures rely on multiple specialist agents coordinated by an orchestration layer. Each agent performs a bounded role. Each produces structured output. Each logs its reasoning and evidence.

This modularity is what makes scalability and auditability possible.

Sutherland’s BFS AI Hub brings together 90+ BFS-specific agents spanning onboarding, underwriting, KYC, AML, transaction monitoring, fraud prevention, disputes and servicing.

Below are three representative agentic workflow patterns that illustrate how autonomy can be applied responsibly in banking.

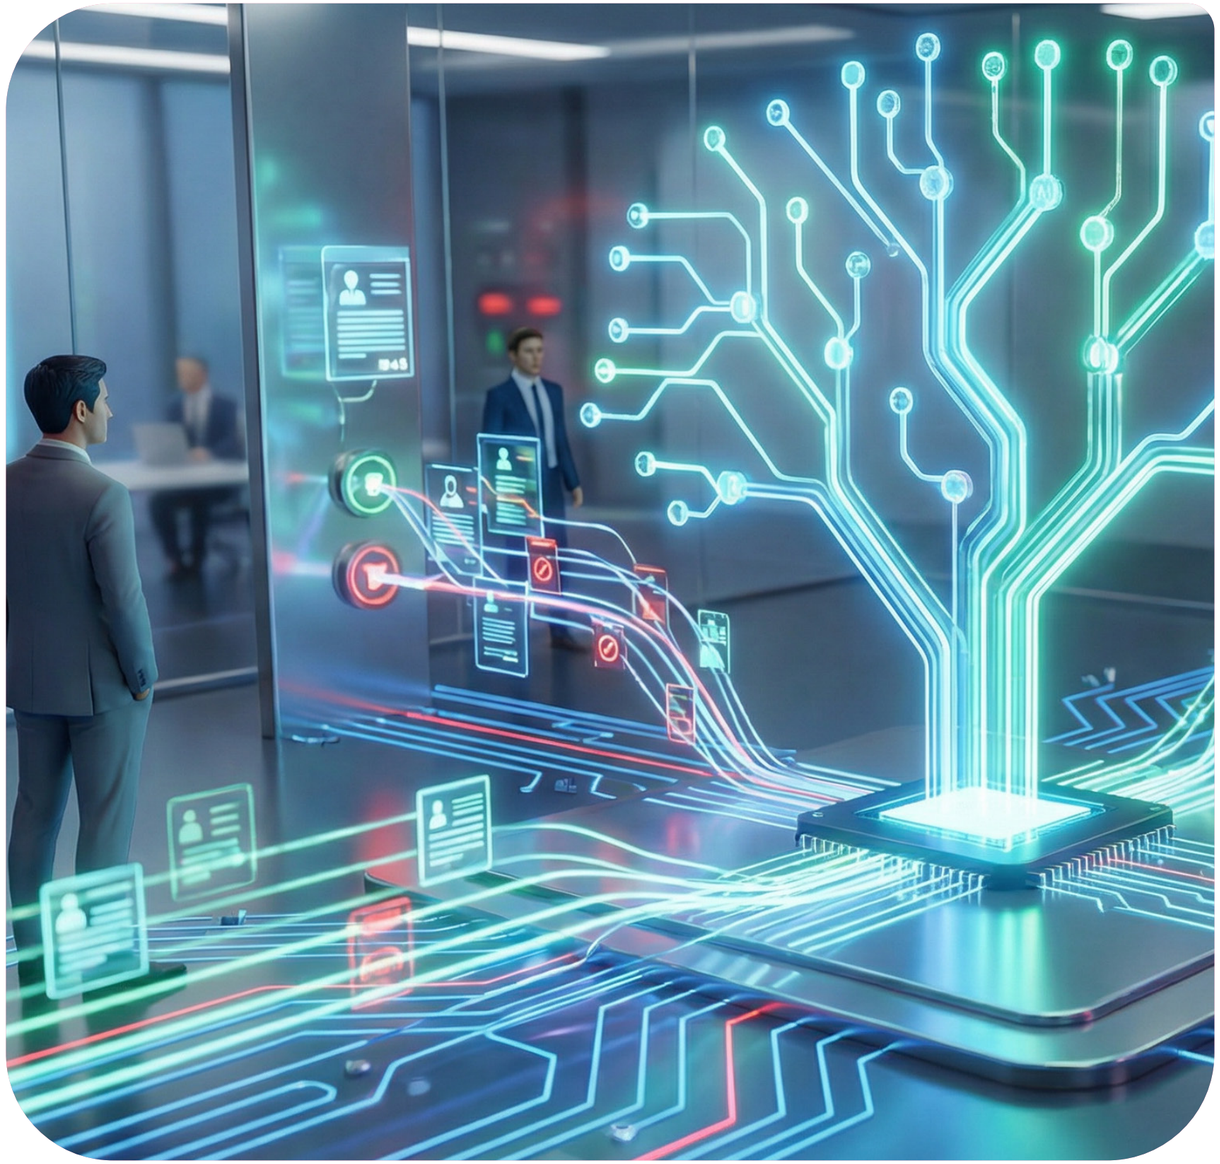

Case pattern 1: **KYC refresh as a governed multiagent workflow**

Periodic KYC refresh is operationally intensive. It involves document retrieval, verification checks, policy validation, customer outreach and compliance documentation.

An agentic KYC workflow typically includes:

- **Ingestion Agent:** Collects customer documents and extracts relevant data fields from structured and unstructured sources.
- **Verification Agent:** Performs identity checks against internal records and approved external databases.
- **Policy Agent:** Applies jurisdictional and risk-tier logic to determine compliance status.
- **Exception Agent:** Flags anomalies and routes complex cases for human review.
- **Documentation Agent:** Generates a structured audit-ready case file including rationale and evidence references.

The orchestrator governs sequencing and escalation. If confidence thresholds fall below predefined limits or new risk signals emerge, the workflow routes to a human analyst.

Why this domain works well for agentic AI:

- Clear policy frameworks
- Structured evidence requirements
- Repeatable lifecycle
- High operational volume

KYC refresh becomes not just faster but more traceable. Each step leaves a digital audit trail that aligns with regulatory expectations.

Case pattern 2: **Disputes as an agentic control loop**

Card disputes and payment chargebacks are governed by strict timelines and evidentiary requirements. They combine structured transaction data with unstructured customer narratives.

An agentic disputes workflow may include:

- **Classification Agent:** Categorizes dispute type based on transaction and narrative inputs.
- **Evidence Aggregation Agent:** Collects transaction logs, merchant data and communication records.
- **Policy Application Agent:** Applies network rules and regulatory requirements.
- **Response Generation Agent:** Drafts customer communications and system updates.
- **Deadline Monitoring Agent:** Tracks regulatory timelines and triggers escalation if thresholds are at risk.

The orchestrator ensures no step is skipped and maintains full traceability.

Why disputes are an ideal proving ground:

- Clear external rulebooks
- High volume
- Measurable economics
- Hard compliance deadlines

If an institution can manage disputes with controlled autonomy and defensible traceability, it builds the governance muscle needed for broader agentic deployment.

Case pattern 3: **Financial Crime Alert Triage as Structured Autonomy**

Financial crime operations are burdened by high alert volumes and significant false positives. Analysts spend substantial time gathering data before applying policy logic.

An agentic alert workflow typically includes:

- **Data Aggregation Agent:** Pulls customer profile, transaction history and risk signals across systems.
- **Contextual Analysis Agent:** Identifies patterns, peer comparisons and anomaly indicators.
- **Policy Assessment Agent:** Applies AML rules and risk thresholds.
- **Disposition Agent:** Recommends closure or escalation based on confidence scoring.
- **Audit Agent:** Generates standardized documentation for compliance review.

Human investigators retain authority over escalations and suspicious activity reporting decisions.

Why this workflow is suitable:

- Highly repetitive data gathering
- Structured policy application
- Clear measurable productivity gains
- Strong need for traceable rationale

When implemented correctly, agentic triage reduces analyst fatigue while strengthening consistency and documentation quality.

In each case:

- The workflow is decomposed into bounded digital roles.
- Orchestration governs sequencing and escalation.
- Human oversight remains embedded.
- Every action generates traceable evidence.

This is not uncontrolled autonomy.
It is structured, policy-aware execution.
That is what defines BFS-grade agentic AI.

8. The economics: building an investable business case

Agentic AI business cases fail when framed as generic productivity. They succeed when framed as unit economics of a process.

Step 1: **Baseline the process:** Quantify current unit cost, cycle time, rework, backlog and risk exposure before introducing agents.

Step 2: **Translate friction into financial impact:** Convert operational delays and exceptions into measurable cost, revenue leakage and compliance risk.

Step 3: **Define autonomy boundaries:** Segment workflow tasks into deterministic, policy-driven and judgment-based activities to control execution scope.

Step 4: **Redesign the workflow:** Reconstruct the end-to-end process around agent execution rather than layering automation onto legacy steps.

Step 5: **Model the autonomy split:** Target a 60–30–10 distribution across autonomous, augmented and human-led work.

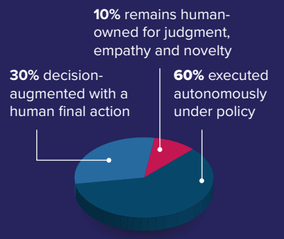

Step 6: **Quantify structural impact:** Measure cost-to-serve reduction, cycle compression, risk mitigation and productivity uplift.

Step 7: **Elevate the narrative:** Position agentic AI as operating model transformation rather than incremental efficiency.

9. Maturity pathway: from pilots to production

Moving from isolated pilots to production is rarely a technology problem alone. It is an organizational and architectural journey that requires clear ownership, measurable outcomes, and an integrated control posture. The maturity pathway shows the sequence in which capability, governance and operating model must evolve so that agentic AI stops being a set of experiments and becomes durable infrastructure. It also clarifies where to invest first so boards and CFOs can see when costs convert into structural value.

Level	Operating reality	What is automated	Governance posture	Outcome
Assist	copilots in teams	summaries, drafting, search	ad hoc	local time savings
Execute	narrow agents	extraction, validation, simple updates	permissions + logging	throughput uplift in a queue
Orchestrate	multiagent workflows	end-to-end case progression	central orchestration + evidence	cycle time reduction with better compliance
Rewire	domain redesign	decisioning + action + learning	control tower integrated with risk	structural cost-to-serve improvement
Autonomous by design	enterprise scale	cross-domain actions with strict guardrails	regulator-ready evidence at scale	differentiated experience with resilience

Table 3. Agentic AI maturity model for BFS

Gartner’s cancellation forecast is effectively a maturity warning label: without governance and value clarity, pilots will die in transition to production.

10. What banking leaders should demand from an AI Hub

If agentic AI is a workforce, an AI Hub is the system that hires, trains, governs and measures that workforce.

From a Gartner lens, four non-negotiables determine production viability:

1. **Orchestration plus tool control**
   Clear action permissions, sequencing, separation of duties and human approvals.

2. **Evidence by default**
   Every decision produces an auditable case file, not a post-hoc narrative.

3. **Industrial testing**
   Scenario testing for accuracy, bias, privacy and resilience before go live and continuously after deployment.

4. **Reusable patterns**
   Standard blueprints for common tasks like document intake, identity checks, policy reasoning and exception handling.

This is exactly why a hub approach matters. It prevents every line of business from reinventing governance and it turns learning into an asset that compounds.

Here’s how Sutherland Agentic AI Hub for Financial Services is architected:

- Foundation

  This covers:

  - **Studio** for building
  - **Orchestrator** for workflow coordination
  - **Trust Center** for responsible AI
  - **Test Center** for real-world validation
  - **Knowledge Center** for curated BFS policy and data
  - **Cloud Center** for secure deployment
- Platforms (Interoperability, Partner ecosystem integration, Monitoring and dashboards
- Agentic Layer (Domain-trained agents, Explainability agents, Workflow orchestrators)

To learn more, visit our **website**

11. Strategic recommendations for Leaders in Financial Services

1. **Treat agents as roles.** Publish boundary conditions and approvals. If you cannot explain what an agent is allowed to do, you cannot govern it.
2. **Start where policy and evidence are strongest:** KYC refresh, disputes and selected servicing workflows. These build traceability muscle fast.
3. **Stand up an AI control tower** spanning technology, risk and compliance to manage agent inventory, approvals, monitoring and value tracking.
4. **Engineer resilience and third-party oversight** using DORA-style discipline even when you are not directly regulated under it.[6]
5. **Assume some projects will fail.** Build faster evaluation, stronger guardrails and clearer outcome metrics early. That is the only way to beat the cancellation curve Gartner is warning about.[2]

2 Gartner, “Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027” (press release, Jun 25 2025).
6 European Banking Authority, “Digital Operational Resilience Act (DORA)” overview.

Conclusion: Why BFS AI Hub is the pragmatic path to scale

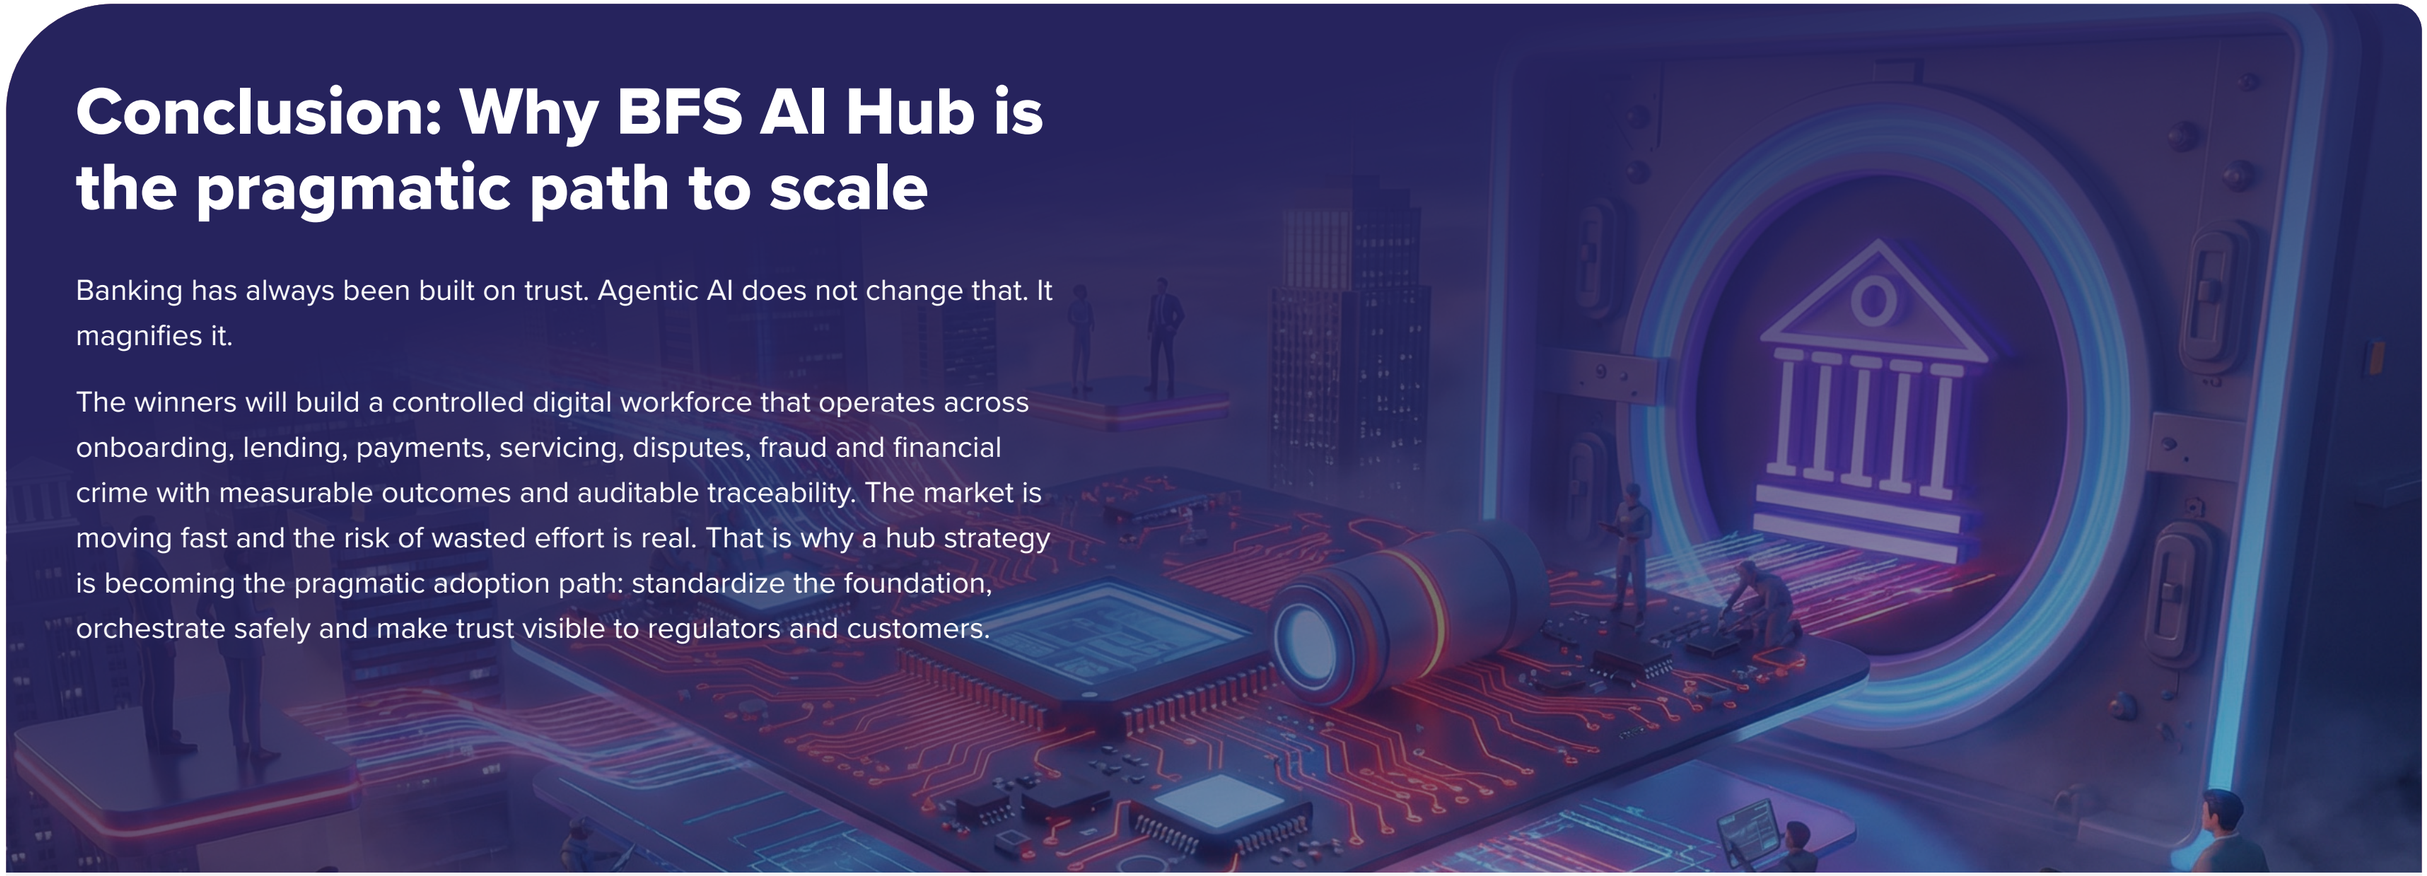

Banking has always been built on trust. Agentic AI does not change that. It magnifies it.

The winners will build a controlled digital workforce that operates across onboarding, lending, payments, servicing, disputes, fraud and financial crime with measurable outcomes and auditable traceability. The market is moving fast and the risk of wasted effort is real. That is why a hub strategy is becoming the pragmatic adoption path: standardize the foundation, orchestrate safely and make trust visible to regulators and customers.

Artificial Intelligence. Automation. Cloud Engineering. Advanced Analytics. For Enterprises, these are key factors of success. For us, they're our core expertise.

We work with global iconic brands. We bring them a unique value proposition through market-leading technologies and business process excellence. At the heart of it all is Digital Engineering Services – the foundation that powers rapid innovation and scalable business transformation.

We've created 363 unique and independent inventions, 250 of which are AI-based and rolled up under several patent grants in critical technologies. Leveraging our advanced products and platforms, we drive digital transformation at scale, optimize critical business operations, reinvent experiences, and pioneer new solutions, all provided through a seamless "as-a-service" model.

For each company, we provide new keys for their businesses, the people they work with, and the customers they serve. With proven strategies and agile execution, we don't just enable change – we engineer digital outcomes.

SUTHERLAND
digital outcomes.

sutherlandglobal.com