


Magpie Projects

COLLECT . CONNECT . CREATE

Agentic AI in Compliance

White Paper

April 2026



Introduction

Since the introduction of large language models (LLMs) to the mainstream public, financial institutions' compliance departments have been brainstorming and testing use cases for groundbreaking technology. Traditional AI models, such as machine learning models, have been in use in compliance departments, e.g. in Anti-Financial Crime (AFC) or Trade Surveillance, far longer already, however LLMs, while impressive, have not yet possessed the tools and capabilities to solve the complex, real-world tasks required to assist in compliance departments such as alert case reviewing or end-to-end investigations involving fragmented data sources, regulatory reasoning, and auditability requirements. Agentic AI frameworks seem to address exactly this hurdle which should allow for many suggested use cases to become technically feasible. Agentic AI refers to tools which utilize agentic frameworks to provide LLMs with autonomy and tool-use capabilities allowing such AI Agents to perform and understand more complex, multi-stepped tasks.

Agentic Frameworks, Opportunities and Limitations

Agentic AI frameworks represent an architectural evolution in artificial intelligence, extending the capabilities of large language models beyond single-turn interactions. Whereas traditional LLMs excel at generating text or predictions in response to discrete prompts, agentic systems are designed to sustain reasoning and action over extended sequences of tasks. They do so by incorporating mechanisms for planning, memory, and external tool use into a coordinated control loop. This loop allows the agent not only to generate language but also to set intermediate objectives, evaluate progress, and adapt its strategy in response to new information. In effect, these frameworks operationalize autonomy, giving AI systems the ability to pursue complex goals that cannot be reduced to a single prompt-response exchange.

The foundations of these systems are becoming increasingly standardized. Planning modules allow an agent to break down high-level objectives into smaller, sequenced steps. Memory components ensure that prior context and decisions are retained, enabling long-running processes rather than isolated exchanges. Tool integrations connect the agent to databases, regulatory feeds, or analytic platforms, turning the language model into a versatile orchestrator of digital workflows. Just as importantly, these frameworks embed mechanisms for collaboration, whether in the form of multi-agent systems that negotiate and divide tasks, or structured channels for human-in-the-loop oversight. Together, these elements transform language models into adaptable problem-solvers capable of navigating dynamic, multi-faceted environments.

The ecosystem of frameworks is expanding rapidly, but they can be analytically distinguished by their architectural emphasis. Some, such as LangChain, Hermes, OpenClaw and LangGraph, prioritize modularity and extensibility, providing researchers

and developers with granular control over memory, planning, and tool integration. Others, including Microsoft's Semantic Kernel and AutoGen, focus on orchestration and enterprise interoperability, embedding agentic functions within broader software ecosystems. Still others, such as CrewAI or OpenAI's Swarm, are exploring distributed or multi-agent paradigms, where coordination and emergent behaviors become the central research problem. From an academic perspective, these approaches can be situated along axes such as centralized versus decentralized control, ephemeral versus persistent memory, and symbolic versus sub-symbolic reasoning support. This diversity of architectures reflects an ongoing search for the most effective abstractions to balance flexibility, reliability, and interpretability.

What is most striking about agentic AI is the range of problems it can be applied to. These systems excel at tasks that require navigating ambiguity, synthesizing diverse information sources, and executing multi-step processes. Unlike traditional machine learning models, which are confined to narrowly defined prediction problems, agents can coordinate across domains, adapt when new information arrives, and pursue goals over time. This makes them uniquely suited to the kinds of dynamic, cross-functional challenges that have historically resisted automation. The vision is of AI that does not simply answer questions, but that participates actively in the work itself: drafting, monitoring, reasoning, coordinating, and ultimately helping organizations act with greater speed and confidence.

Of course, technology is not without its boundaries. Agents can still inherit the weaknesses of the models they are built on, including occasional errors or shallow reasoning. Integration with external tools can introduce points of fragility, and transparency remains a work in progress. Yet these limitations should be viewed as milestones along a trajectory of rapid improvement, not as hard barriers. Each generation of frameworks is addressing these issues directly, with advances in reasoning methods, auditability, and guardrail design already narrowing the gap between what is possible in theory and what is reliable in practice.

Looking forward, the research trajectory suggests increasing convergence between theoretical AI planning and practical enterprise deployment. Early implementations may appear brittle, but the pace of iteration indicates that core challenges such as explainability, alignment, and scalability will continue to see rapid progress. In regulated sectors such as finance, this evolution will require a careful balance: frameworks must deliver technical robustness while simultaneously providing the transparency and auditability that compliance demands.

High Impact Areas for Compliance

The potential for agentic AI to reshape compliance becomes most evident when examining the structural characteristics of compliance work. Across functions, many activities are defined by recurring workflows that require synthesizing information from multiple sources, applying rule-based or policy-driven logic, and producing outputs that are fully documented and auditable. These attributes align closely with the strengths of agentic frameworks, which combine persistence, reasoning, and structured execution.

Rather than being tied to specific functions, the applicability of agentic AI can be understood through a set of recurring patterns. High-impact opportunities typically share three defining characteristics. First, they involve multi-step processes that require coordination across systems, such as gathering data, enriching context, and applying decision logic. Second, they are high-volume and repetitive, creating significant operational cost when performed manually. Third, they require traceability and explainability, as outputs must be defensible to internal stakeholders and regulators.

Activities that combine these traits tend to create bottlenecks within compliance organizations. They are often too complex to be fully automated using traditional rule-based systems, yet too structured to require continuous human intervention. As a result, they sit in an operational “middle ground” where human effort is heavily relied upon for orchestration, rather than judgment.

Agentic AI is particularly well suited to addressing this gap. By decomposing workflows into manageable steps, maintaining context across tasks, and documenting intermediate reasoning, agentic systems can execute these processes in a more consistent and scalable manner. Importantly, this does not eliminate the role of compliance professionals, but instead shifts their focus toward oversight, interpretation, and decision-making in higher-risk scenarios.

The following use cases illustrate how these characteristics manifest in specific compliance functions and how agentic AI can be integrated into existing workflows to deliver tangible improvements in efficiency, accuracy, and responsiveness.

Specific Compliance Use Cases for Agentic AI

The high-level promise of agentic frameworks becomes most tangible when applied to well-defined compliance functions that are both cost-intensive and operationally critical. In the following use cases, we illustrate how agentic AI can be integrated into existing workflows to enhance efficiency, accuracy, and responsiveness. While the underlying goals of these functions remain unchanged, agentic systems provide the scaffolding for more structured, persistent, and auditable execution.

Gap Analysis of Compliance with New and Changing Regulations

Compliance departments continually face the challenge of assessing how evolving regulations affect existing policies, procedures, and controls. The current workflow typically requires regulatory specialists to monitor updates from regulators, interpret the language of new rules, and manually cross-reference them against internal policy frameworks. This process is slow, error-prone, and heavily dependent on human expertise, often leading to significant delays in remediation and increased exposure to regulatory risk.

An agentic AI solution can transform this process by automating the initial stages of regulatory mapping and comparison. An agent could continuously monitor regulatory feeds, extract the relevant provisions, and align them with internal policy documents using semantic search and clause-level analysis. By maintaining a persistent memory of prior regulatory mappings, the agent could also identify recurring gaps across

jurisdictions or business units. Once discrepancies are detected, the system could generate structured reports, highlighting specific clauses requiring revision and even suggesting draft language based on established templates, all while citing the lines in the regulations causing the changes for easier review. Human compliance officers would remain central to interpretation and approval, but their workload would be shifted from exhaustive manual review toward targeted oversight. This integration reduces the time between regulation issuance and institutional alignment, minimizing both compliance costs and regulatory risk exposure.

Insider Trading Detection and Investigation

Detecting insider trading remains one of the most complex challenges in compliance. Traditional surveillance systems rely on static rules and anomaly detection algorithms, which generate large volumes of alerts but often lack the contextual depth needed to evaluate suspicious behavior. Investigations demand piecing together data from multiple systems: trading records, communications archives, employee access logs and then reconstructing timelines of activity. This is a resource-intensive task requiring skilled investigators and lengthy turnaround times.

Agentic AI introduces a more dynamic approach to both detection and investigation. Once an alert is generated by a monitoring system, an agent could autonomously collect contextual information across systems: retrieving relevant communications (emails, chat logs, voice transcripts), correlating these with employee access to material non-public information (e.g., financial results, deal documents), and aligning them with executed trades. The agent's planning and reasoning capabilities allow it to construct candidate timelines of events, identify potential conflicts of interest, and flag inconsistencies or patterns indicative of insider activity. Importantly, the system would not replace human investigators but instead act as an investigative assistant, producing structured dossiers that can be rapidly validated or dismissed. By reducing the manual effort of data gathering and synthesis, agentic systems free compliance experts to focus on higher-order judgment and case escalation.

Market Event Materiality Assessment

Another critical function in compliance is the assessment of whether a market event, such as a significant corporate development or external crisis, constitutes "material information" under securities regulations. Currently, this assessment is highly manual, requiring compliance teams to monitor news feeds, evaluate relevance to institutional holdings, and make time-sensitive judgments on disclosure and trading restrictions. Delays in recognizing materiality can lead to regulatory breaches or reputational damage.

Agentic AI can play a role in both monitoring and assessment. Agents can continuously scan structured and unstructured data sources, including financial news, press releases, social media, and internal research outputs. By applying reasoning chains, an agent could match emerging events against predefined thresholds or contextual patterns, such as revenue impact, sector exposure, or regulatory disclosure requirements. When a candidate material event is detected, the system could initiate a structured assessment workflow: drafting a preliminary analysis, cross-referencing relevant regulatory criteria, and routing the case for expedited human review. With persistent memory, the agent

could also track prior assessments, allowing compliance officers to benchmark decisions against historical precedents. This capability enhances both the speed and consistency of materiality determinations, reducing the risk of oversight in volatile market conditions.

Name List Screening

Screening client, counterparty, and transaction data against regulatory and internal watchlists is one of the most resource-intensive compliance activities. Current systems rely on rule-based matching techniques that generate significant volumes of false positives, requiring manual review to confirm matches and clear false alerts. The scale of this activity, particularly in large financial institutions, translates into substantial compliance cost.

Agentic AI offers a pathway to reduce these inefficiencies by augmenting traditional screening with contextual reasoning. An agent could manage the screening workflow end-to-end: parsing new client or transaction data, applying fuzzy matching algorithms against multiple watchlists, and then gathering contextual information (e.g., geographic details, corporate affiliations, prior screening outcomes) to evaluate the likelihood of a true match. By iterating across multiple data points and maintaining memory of past screening decisions, the agent can progressively reduce false positives and surface the cases most likely to require escalation. Moreover, because agents can document their reasoning steps, the screening process becomes more transparent and auditable, addressing regulatory concerns about explainability. While ultimate decision-making would remain with compliance officers, the agent's ability to integrate, reason, and prioritize fundamentally reshapes the efficiency of screening operations.

Conclusion

Agentic AI represents a pivotal shift in the technological foundations of compliance. By extending large language models with planning, memory, and tool-use capabilities, these frameworks enable systems to tackle the multi-step, context-rich tasks that define modern regulatory environments. As demonstrated across areas such as regulatory change management, insider trading detection, market event assessment, and name list screening, agentic solutions can both relieve the operational burdens faced by compliance teams and enhance the quality, consistency, and auditability of their outputs.

At Magpie Projects, our guiding principle is simple: Collect. Connect. Create. This ethos applies directly to the future of compliance: collecting the right data, connecting it across fragmented systems, and creating insights that drive both efficiency and resilience. We invite institutions, regulators, and partners to join us in this transformation. For further information, or to discuss how agentic AI can be tailored to your compliance operations, please reach out to continue the conversation. For financial institutions, the imperative is no longer whether to explore these systems, but how to integrate them thoughtfully.

Authors



Philipp Faulstich, CFA

Senior Manager

M: +49 151 25141204

E: philipp.faulstich@magpieprojects.com




Benjamin Cohen Harel

Senior Consultant

M: +49 175 8869251

E: benjamin.cohenharel@magpieprojects.com



Emerging from LPA Consulting, we bring over two decades of capital markets expertise but now sharpen our focus on forward-looking strategic transformation.

We act on early insights to help clients stay ahead, transforming intelligence into actionable strategies and guiding clients from traditional capital markets toward the more integrated financial markets, where public and private capital flows intersect.

We stay at the forefront of change, helping clients navigate shifts—whether in policy, technology, or changing structures. We ground ambition in experience, enabling faster decisions and measurable outcomes. As financial markets evolve, so must the tools. We embed AI and use your data to build smarter, scalable systems, always applying technology with purpose.