

**OLIVER WYMAN**  
A MARSH BUSINESS

# Reimagining Compliance With Agentic AI

Enabling Proactive, Continuous Compliance  
in a Complex Regulatory World



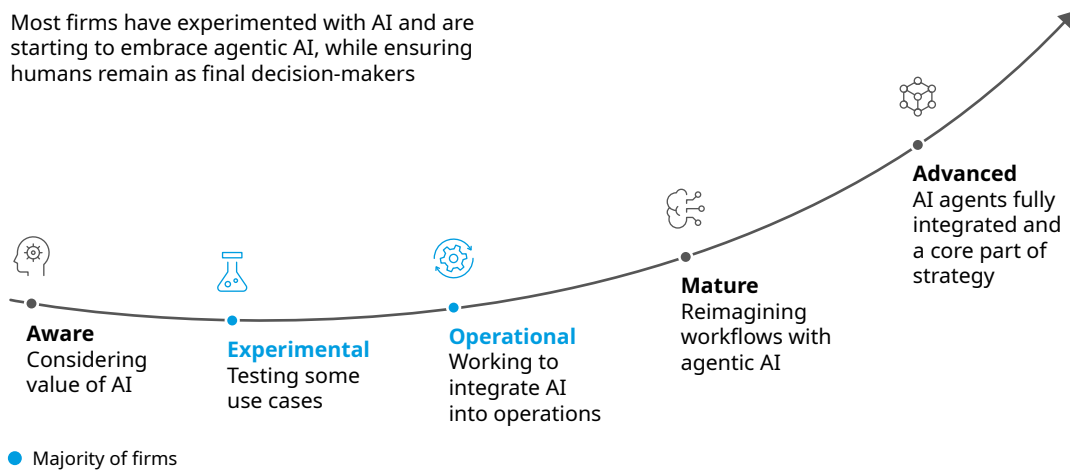
David Choi  
David Carretero  
Allen Meyer  
Daniel Tannebaum  
Frederic Thomas-Dupuis  
Kim-Mai Le

Compliance teams should ambitiously embrace agentic AI systems that are capable of performing semi-autonomous tasks to reimagine processes, unlock greater efficiency, and enable proactive, continuous risk management.

Compliance leaders today face unprecedented challenges, including diverging regulatory expectations, mounting cost pressures, and the urgent need to enhance operational effectiveness. AI presents a significant opportunity to reshape how compliance operates, yet our benchmarks and experience indicate that most organizations remain in early stages — experimenting with generative AI in narrow, incremental ways rather than reimaging workflows end-to-end. While helpful, these early pilots often fail to achieve the transformative impact and return on investment anticipated.

### Exhibit 1: Firms experimenting with AI

Most firms have experimented with AI and are starting to embrace agentic AI, while ensuring humans remain as final decision-makers



Source: Oliver Wyman analysis

True transformation requires more than sprinkling AI into today’s workflows. The next wave, now emerging among early AI leaders, demands reimagining the compliance function end-to-end: how data is used, how work is divided between people and machines, and how compliance teams deliver value to business leadership and regulators.

Compliance teams must think bigger and embrace agentic AI, semi-autonomous systems (“agents”) that can plan, act, and coordinate across multi-step workflows with human oversight.

#### Agentic AI

Semi-autonomous systems that can plan, act, and coordinate across multi-step workflows with human oversight, moving from “AI as a tool” to “AI as a co-worker.”

Operating within explicit guardrails, each AI agent is designed to perform a specific task — such as monitoring for regulatory changes, performing due diligence checks, or preparing case documentation — and escalate judgment calls to humans. To enable trust and confidence, AI outputs should be transparent, demonstrating data sources, lineage and rationales, and reproducible runs.

Our benchmarks and client engagements have shown that reimagined AI-enabled compliance functions can enhance output quality and risk detection by up to four times, enable 100% continuous and comprehensive testing and coverage, and reduce time spent on manual, repetitive tasks by 50 to 70%, freeing capacity for higher-value work. Agentic AI enables continuous data collection, analysis, and insight generation across complex compliance workflows, dramatically enhancing efficiency, accuracy, and scalability. By flipping the paradigm from “expert-created” to “expert-reviewed,” compliance teams can streamline operational workflow, empower staff to focus on judgment and strategy, and achieve proactive, continuous compliance.

That said, realizing these benefits will require careful governance, risk management, and planning. Alongside the significant potential gains, organizations should anticipate new development and operational costs, such as software build, integration, validation, and model maintenance. Agentic AI also introduces new forms of model risk, explainability challenges, and operational dependencies. Finally, achieving the transformation will require a clear roadmap: defining the ambition and measurable objectives, redesigning human-AI roles, and executing a disciplined rollout supported by strong governance, testing, and workforce enablement.

## **Agentic AI unlocks continuous, intelligent compliance operations that detect and manage risks more effectively**


Agentic AI enables a fundamental shift in how compliance functions operate. AI agents operate “always on,” enabling near real-time identification and escalation of risks.

For example, agentic AI can autonomously track regulatory changes and policy updates, ensuring compliance teams stay ahead of evolving requirements. Similarly, in Know Your Customer (KYC) and customer due diligence (CDD) processes, agentic AI enables dynamic, ongoing monitoring of customer risk profiles, allowing timely adjustments to due diligence based on new data.

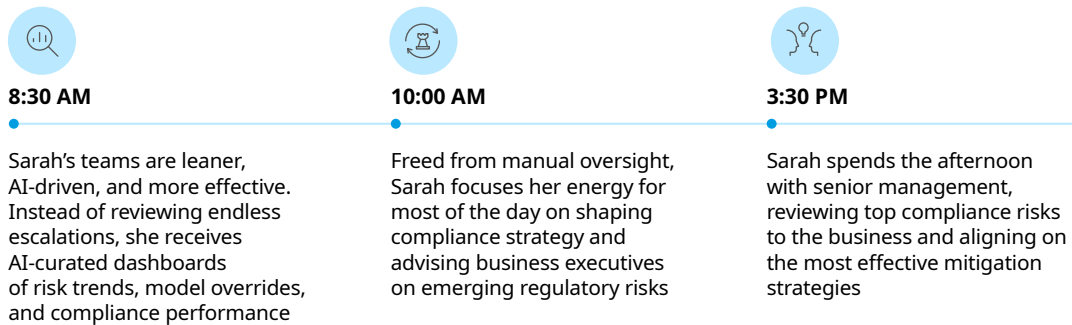
The art of the possible extends even further. In the long run, AI has the potential to upend the entire paradigm of compliance monitoring. Today, banks rely on rules-based alerts because transaction volumes are too great to review exhaustively. With near-unlimited AI capacity, firms could one day analyze 100% of transactions in real time — not just flagging issues after the fact but stopping risky transactions before they occur.

## Exhibit 2: A day in the life of teams in 2030 reimagined with AI

### Investigations AI operations manager



### Compliance adviser

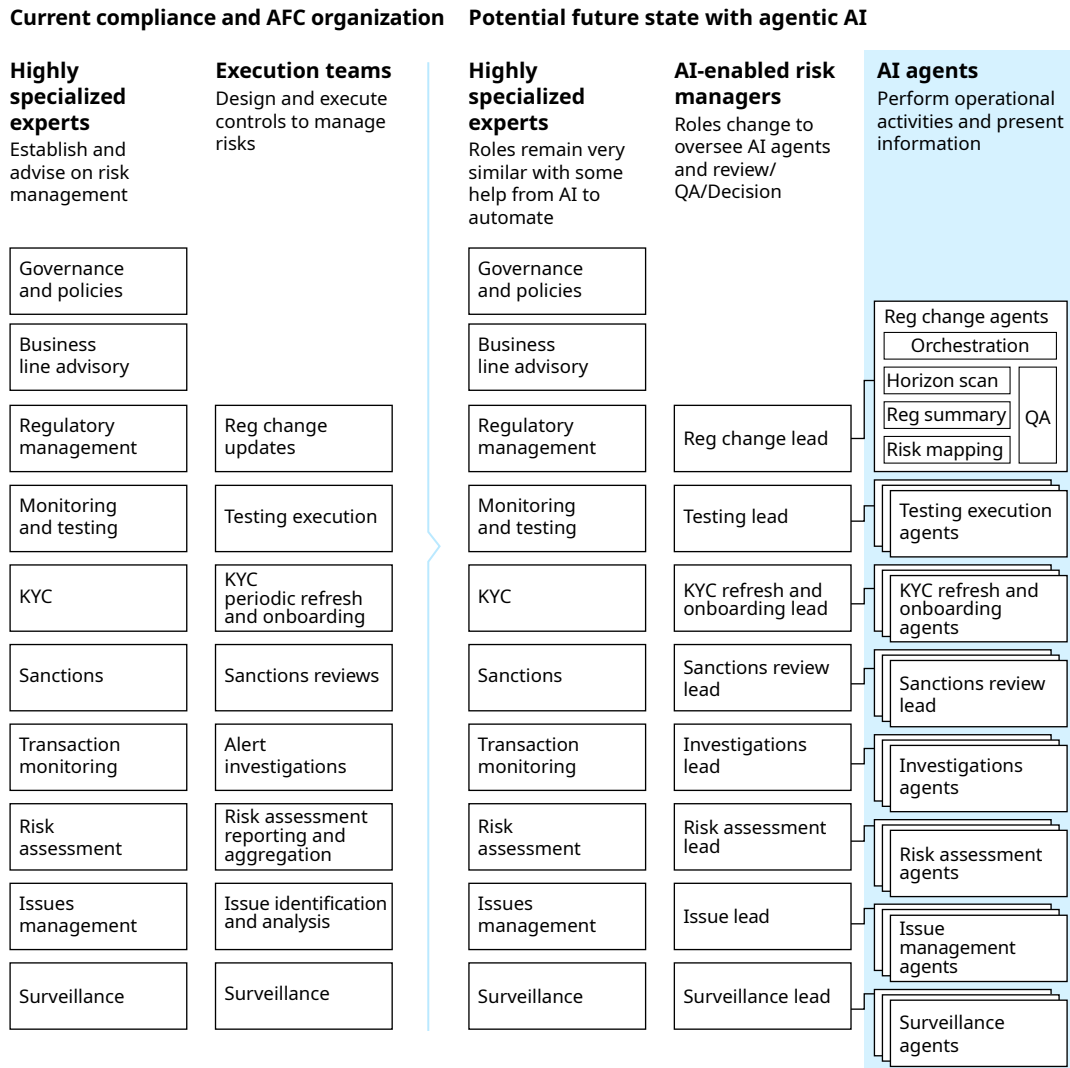


Source: Oliver Wyman analysis

## How the compliance workforce will change in the future

Today, most compliance functions remain headcount-heavy and operationally skewed. Highly trained staff spend a small share of their time on specialized work such as governance, advisory, and risk assessments, compared with execution tasks such as investigating cases, testing controls, reconciling data, and processing regulatory changes. These activities are essential for regulatory compliance but do little to elevate the strategic value of compliance and take away time that could better be spent on risk decisioning.

### Exhibit 3: How a compliance and anti-financial crime team might evolve with the use of agentic AI



Notes: AI agents included are illustrative and non-exhaustive  
Source: Oliver Wyman analysis

The future looks radically different. With agentic AI, compliance functions can fundamentally reshape their operating model. The need for highly specialized subject experts will remain but execution teams will be materially recast as AI-enabled risk managers who configure agentic playbooks, supervise agentic performance, and retain final decision-making and accountability rather than overseeing large pools of manual operators. Routine, operational activities will be executed end-to-end by AI agents and automatically quality-checked by AI agents, elevating human talent to focus on high-value work like final adjudications, exceptions handling, and targeted QA where stakes are highest. In practice, execution leads will oversee teams of AI agents and review/approve their outputs, ensuring faster cycle times, fewer manual handoffs, and more defensible outcomes without relinquishing control.

Instead of drowning in alerts and paperwork, compliance experts will spend their time interrogating AI insights, weighing trade-offs, and engaging with business leaders and regulators on forward-looking strategy and risk management. This shift unlocks the real potential of compliance talent. Analysts become decision-makers supported by rich, real-time intelligence. Advisory teams step out of reactive firefighting and into proactive guidance. Leaders gain the headspace to shape compliance strategy as a competitive advantage. In other words, far from replacing people, agentic AI elevates them.

## Agentic AI will improve efficiency and ultimately lead to better risk management

Our benchmarks indicate that some compliance processes, such as testing, require 50% to 70% of time on manual, repetitive tasks that could be automated by agentic AI. The efficiency gains and improved risk coverage translate into substantial cost savings and enhanced risk management.

### Exhibit 4: Example risk management improvements and impacts achieved with AI

Improvements	Description	Examples	Example impacts achieved
High-quality outputs and better risk decisions	Improved accuracy and decisioning in risk evaluations	<ul style="list-style-type: none"> <li>• Sharper sanctions dispositions</li> <li>• Better fraud and AML prevention</li> </ul>	<b>2-4x</b> more risk identified in transaction monitoring and fraud
	Richer, more consistent, auditable documentation	<ul style="list-style-type: none"> <li>• Detailed SARs that enable regulators and law enforcement</li> <li>• Plain English summaries of laws, rules, and regulations that are easier to interpret</li> </ul>	
Continuous and comprehensive testing and assurance	100% risk management review and testing rather than sampling	<ul style="list-style-type: none"> <li>• Continuous QA of all CDD files</li> <li>• Full testing of controls inventories</li> </ul>	<b>100%</b> full and continuous QA and controls testing
	Real-time, ongoing, and proactive monitoring of risk factors	<ul style="list-style-type: none"> <li>• Perpetual KYC refreshes</li> <li>• Always-on regulatory change tracking and updates</li> </ul>	
Less manual work, more time for complex tasks	Time and cost saving in more efficient risk management workflows	<ul style="list-style-type: none"> <li>• Reduced SAR turnaround times</li> <li>• Cost savings from controls optimization and automation</li> </ul>	<b>50%-75%</b> reduction in investigation times
	Streamlined and automated processes to unlock time for judgement-based risk activities	<ul style="list-style-type: none"> <li>• Collapsing L1/L2 model in investigations</li> <li>• Reduce number of manual steps</li> </ul>	
Improved customer and employee satisfaction	Teams focus on high-value tasks, not drudgery	<ul style="list-style-type: none"> <li>• Investigators reviewing only high-risk alerts</li> <li>• Experts spending time on escalations</li> </ul>	<b>30%-50%+</b> more accurate customer risk ID with fewer outreach requests
	Smoother customer interactions and increased trust	<ul style="list-style-type: none"> <li>• Faster customer onboarding</li> <li>• Fewer KYC refreshes on unchanged profiles</li> </ul>	

Notes: Estimated impacts are based on Oliver Wyman client projects and benchmark analyses

## Clearly defined autonomy bounds for agentic AI ensure clarity, accountability, and compliance

Implementing agentic AI also introduces new inherent risks, ranging from model bias and hallucinations to lack of transparency or insufficient human oversight. These risks could erode regulatory trust or create new vulnerabilities. To mitigate them, agentic AI should be structured with clearly defined autonomy bounds. Each AI agent is purpose-built and designed for specific tasks with explicit guardrails and escalation criteria — for example, a planner agent that sequences steps, a QA agent that validates outputs, or a data-gathering agent that assembles evidence. Each agent operates within pre-defined playbooks and hands-off judgment calls at defined thresholds for humans to review and make decisions.

To reinforce trust and accountability, the system should be coordinated by an orchestration layer with observability (such as trace logs, lineage, audit trails) and operate with explicit controls embedded in each domain that clearly show how information is provided and used as rationale for decision-making.

Example controls for AI agents for regulatory change management:

- Source citations must reference the law/rule/regulation and pinpoint the location (section, clause, paragraph, sentence) used to derive each requirement
- Change differences must show redline comparisons between prior and updated regulatory texts and between prior and updated internal policy/process wording, with anchors back to the specific regulatory clauses driving each change

Example controls for AI agents for KYC:

- Risk scoring rationale must show which features/signals were used (such as geography, occupation, adverse media hit types), the weighting/logic, and the resulting score with thresholds and reason codes
- Adverse media logging must capture article URLs, publishers, publication dates, entity disambiguation steps, language translation notes, and relevance grading, with links stored in the case record

## Compliance leaders can realize this transformation at scale by adopting AI carefully and strategically

A structured and well-planned implementation ensures that AI deployment aligns with regulatory requirements and organizational goals, minimizing risks while maximizing benefits. By prioritizing a measured approach, compliance teams can harness AI's power to enhance efficiency, improve decision-making, and strengthen risk management across the enterprise — paving the way for sustainable, scalable success.

### 1. Frame the ambition and prioritize

---

**Set a clear ambition and measurable objectives to align expectations.** Define a clear vision that communicates the importance and ambition of this transformation to the organization (for example, “We will re-engineer three to five end-to-end processes with agentic AI in the next 12 months”). Articulate the scope, including guardrails and what will remain human judgment by design, and align on measurable targets (such as cycle-time, quality, coverage).

**Select the right opportunities to start with.** Rank candidate processes by impact (such as FTE hours, manual effort, and error rates) and feasibility (such as data availability/quality and timing/risk/regulatory sensitivity). Select two to three high ROI areas enabling early value-creation and credible learning without undue risk.

### 2. Reimagine the process and roles of humans and AI

---

**For each priority, reinvent the workflow and AI/human roles to achieve desired outcomes.** Revisit the outcome the process is meant to achieve, then redesign the end-to-end workflow considering what both agentic AI and humans should do. Define the agentic playbook, including agent mandates, autonomy bounds, escalation thresholds, and evidence required, so every action is explainable and reviewable.

### 3. Develop a disciplined rollout strategy and controls to mitigate AI risks

---

**Execute a risk-managed rollout using quality gates.** Run agents in parallel against BAU processes to compare outcomes and decisions. Use test packs, challenger scenarios, and reproducibility checks to validate performance. Promote to production only when quality gates are met.

**Build workforce capability and accountability to operate at scale.** Define the target roles and competencies required. Reskill today's operators into tomorrow's AI operations managers, shifting capacity from low-value tasks to judgment-intensive activities.

Over time, compliance teams should institutionalize the approach above by establishing an AI operations factory that standardizes playbook templates, validation protocols, and release governance and continuously enables a portfolio of reinvention that self-funds with efficiency gains.

## Agentic AI is the future of compliance

Compliance leaders are under pressure to deliver more with less, and agentic AI is the best transformative path forward. By reimagining compliance processes with semi-autonomous AI agents that amplify human expertise, organizations can unlock significant ROI, shift from reactive to proactive risk management, and elevate compliance teams into strategic advisers. This is not about eliminating people, but elevating people — freeing experts from manual drudgery to become trusted voices to the business and regulators.

The time to act is now: Firms that embrace agentic AI today will be leaders in future-proofing compliance and meeting the challenges of an increasingly complex regulatory landscape.

Oliver Wyman, a Marsh business, is a management consulting firm combining deep industry knowledge with specialized expertise to help clients optimize their business, improve operations and accelerate performance. Marsh is a global leader in risk, strategy and people, advising clients in 130 countries across four businesses: Marsh, Guy Carpenter, Mercer and Oliver Wyman. With annual revenue of \$24 billion and more than 90,000 colleagues, Marsh helps build the confidence to thrive through the power of perspective.

For more information, visit [oliverwyman.com](http://oliverwyman.com), or follow on LinkedIn and X.

Americas  
+1 212 541 8100

Europe  
+44 20 7333 8333

Asia Pacific  
+65 6510 9700

India, Middle East & Africa  
+971 (0) 4 425 7000

#### AUTHORS

**David Choi**  
Partner

**David Carretero**  
Partner

**Allen Meyer**  
Partner

**Daniel Tannebaum**  
Partner

**Frederic Thomas Dupuis**  
Partner

**Kim-Mai Le**  
Principal

Copyright ©2026 Oliver Wyman

All rights reserved. This report may not be reproduced or redistributed, in whole or in part, without the written permission of Oliver Wyman and Oliver Wyman accepts no liability whatsoever for the actions of third parties in this respect.

The information and opinions in this report were prepared by Oliver Wyman. This report is not investment advice and should not be relied on for such advice or as a substitute for consultation with professional accountants, tax, legal or financial advisors. Oliver Wyman has made every effort to use reliable, up-to-date and comprehensive information and analysis, but all information is provided without warranty of any kind, express or implied. Oliver Wyman disclaims any responsibility to update the information or conclusions in this report. Oliver Wyman accepts no liability for any loss arising from any action taken or refrained from as a result of information contained in this report or any reports or sources of information referred to herein, or for any consequential, special or similar damages even if advised of the possibility of such damages. The report is not an offer to buy or sell securities or a solicitation of an offer to buy or sell securities. This report may not be sold without the written consent of Oliver Wyman.