Thomson Reuters™

How to use agentic AI workflows in professional services

An introduction to agentic AI use cases, strategies, and benefits for legal, tax, and risk professionals

Contents

Understanding the AI landscape

Professional services industries, including those home to specialized knowledge such as legal, tax, and risk organizations, are at a pivotal crossroads when it comes to artificial intelligence (AI). Once familiar only to elite technology companies, AI has officially entered the mainstream and is proving useful across a spectrum of roles and industries. According to Gartner, more than 80% of enterprises will have used generative AI APIs or deployed generative AI–enabled applications by 2026.[1] While this movement into everyday workflows began with generative AI, it is agentic AI that's now reshaping how professionals conduct tasks—and its presence is growing.

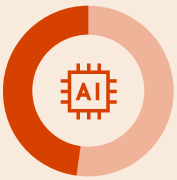

48%
of tech executives are already adopting agentic AI.[2]

50%
of tech executives expect that more than half of their AI deployments will be autonomous within the next two years.[2]

78%
of organizations reported using AI in 2024, up from 55% the year before.[3]

53%
of professionals say their organizations are already experiencing ROI from AI adoption.[4]

Common misconceptions

The momentum is building and AI is here to stay, but many professionals remain cautious. Agentic AI is often misunderstood as an advanced chatbot, or a replacement for generative AI, but it often works in combination with other technology, including generative AI, to plan and execute work. While generative AI responds to prompts by producing content such as text or images, agentic AI goes further: it can make decisions, take actions, and adapt to changing environments and information with minimal human input. Unlike chatbots, which follow predefined scripts and handle simple tasks, agentic AI systems can execute multi-step workflows—for example, researching a legal issue, drafting an insurance document, or scheduling client appointments for the upcoming tax season. Unlike generative AI, agentic AI can break down goals into subtasks, execute them, evaluate outcomes, and adapt based on results.

This human-like ability has fostered the misconception that agentic AI can replace professionals, when in fact, it still requires human oversight, interpretation, and validation. It cannot conduct nuanced conversations, exercise ethical judgment, or fully understand context without oversight—important characteristics of successful legal, tax, and risk professionals. Instead, it complements human expertise by streamlining repetitive tasks and enabling professionals to focus on strategic, high-value work. In short, it is a powerful collaborator that helps people move faster and make more room for creativity.

Want a deeper dive into how agentic AI works and where it's headed? Read the ebook: Agentic AI 101.

1. "Gartner Says More Than 80% of Enterprises Will Have Used Generative AI APIs or Deployed Generative-AI Enabled Applications by 2026," Gartner, 2023.
2. Technology Pulse Poll, Ernst & Young LLP, 2025.
3. The 2025 AI Index Report, Human-Centered Artificial Intelligence, Stanford University.
4. Future of Professionals Report, Thomson Reuters, 2025.

The benefits of agentic AI in professional workflows

Professional-grade agentic AI unifies disparate information sources and integrates with internal and external systems such as company files, databases, trusted external content like Westlaw, and popular tools like Microsoft 365. This allows people to surface insights faster, verify results with confidence, and make better-informed decisions—without switching between disconnected tools or risking incomplete context. The benefits are horizontal (company-wide) and vertical (industry- and role-specific), and include the ability to:

Accelerate complex workflows: Agentic, guided workflows help complete multi-step, transactional tasks more efficiently.

A global tax team leverages agentic AI to categorize over tens of millions of transactions for sales and use tax compliance. The AI not only classifies the data, but also flags anomalies, updates tax codes, and generates audit-ready reports without manual intervention.

Access trusted content: When grounded in authoritative industry sources, agentic AI can check that AI-generated insights are verifiable and legally sound.

Boost productivity: By uniting research, analysis, and drafting in one platform, and not having to switch between tools, professionals can handle more tasks and also reallocate time for creative, human-centric work.

A financial risk team uses agentic AI to investigate suspicious transactions faster and more accurately. The AI agent surfaces highly relevant insights related to the party and counter parties to help close out investigations quicker, freeing analysts to focus on strategic risk mitigation.

Enhance quality and accuracy: Agentic AI solutions that include expert-designed question and task ("prompt") libraries help you get reliable results without the time-consuming trial and error of manual prompt tuning.

A law firm uses agentic AI to simultaneously analyze internal case files and external legal databases. When preparing a motion, the AI cross-references prior firm cases with current case law to draft a first version of the brief, saving hours of manual research and ensuring consistency and accuracy.

53%
believe their organization is already experiencing at least one type of benefit from AI adoption*

Anticipated individual savings due to AI
5 hours per week
240 hours per year*

Around half expect to have begun **their transformation journey** in next 12 months*

* Future of Professionals Report, Thomson Reuters, 2025.

Navigating barriers to adoption

Despite the promise of agentic AI, some legal, tax, and risk professionals remain hesitant to adopt it. This paralysis is rooted in uncertainty, complexity, and a fear of falling behind without a clear path forward, leaving late adopters to fall behind exponentially.

Disconnect between future aspiration and current pace

80%
believe that AI will have a high or even transformational impact on their profession within 5 years,* but...

only 38%
expect to see transformational or high levels of change in their firm this year.*

53%
believe their firm is already experiencing at least one type of benefit from AI adoption,* but...

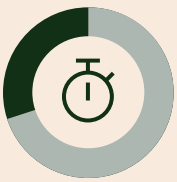

30%
believe their firm is moving too slowly in AI adoption.*

* Future of Professionals Report, Thomson Reuters, 2025.

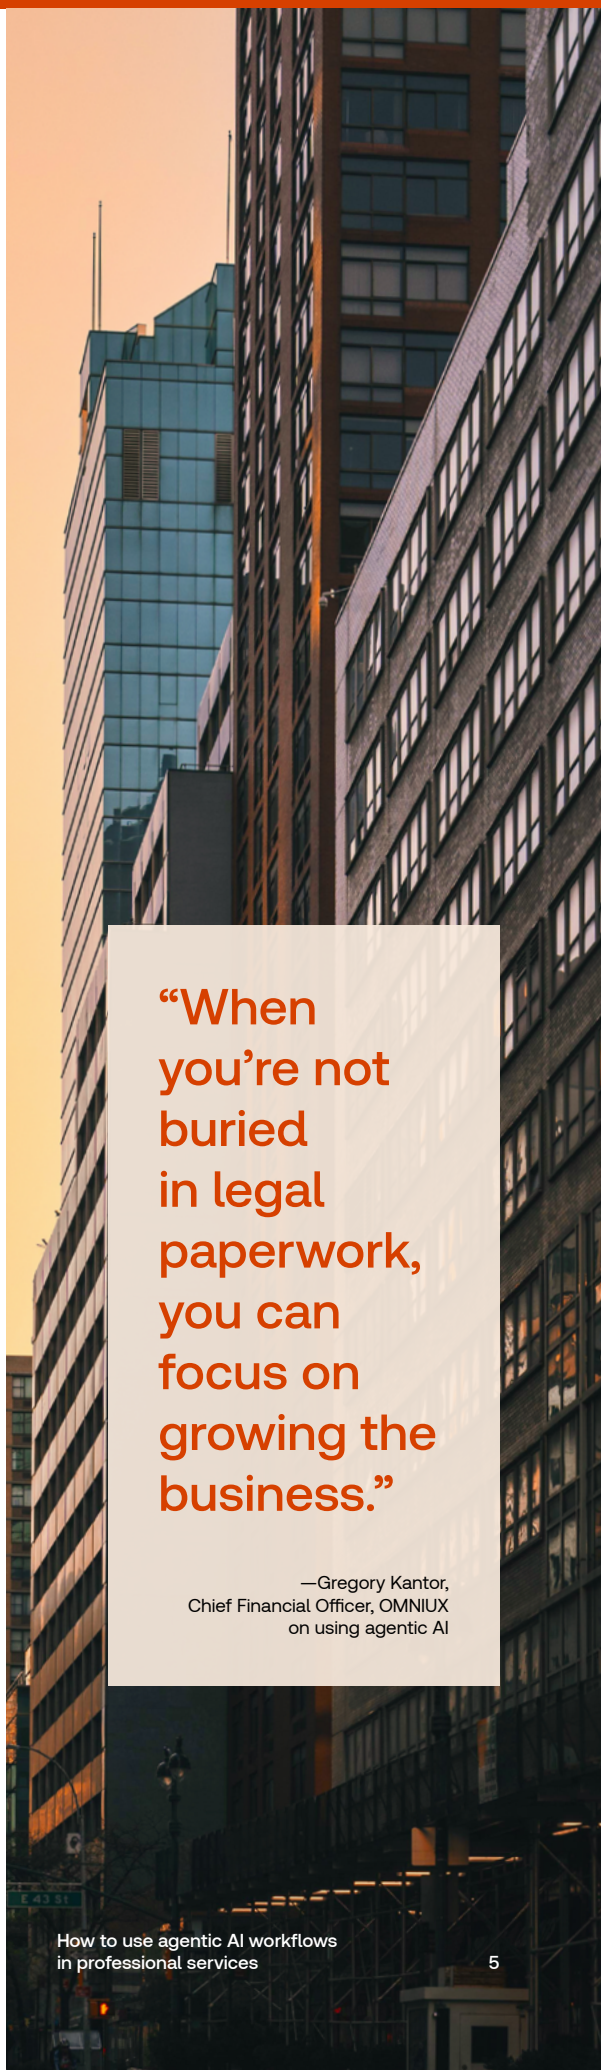

> "When you're not buried in legal paperwork, you can focus on growing the business."
>
> —Gregory Kantor, Chief Financial Officer, OMNIUX on using agentic AI

Accuracy, budget, and security: perceived barriers to investment

Negative consequences of most concern

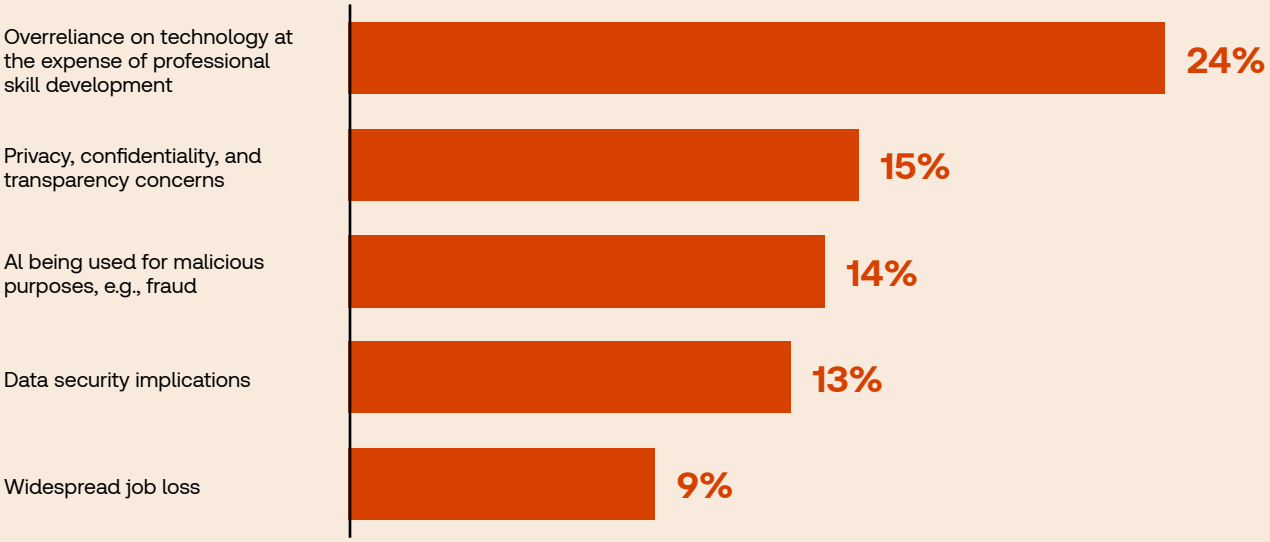

Barriers to investment

The big five: barriers to agentic AI adoption

Knowing your team's roadblocks is the first step to overcoming them

1 Paralysis by complexity

Organizations that are outside of the early adopter tier—those in "step zero"—are overwhelmed by the pace of change. They've heard the buzz around generative AI and agentic AI, but don't know where to begin. Concerns about billable hours, client confidentiality, and talent disruption often cloud the conversation, leading to inaction. Those that experience this paralysis by complexity can move forward by understanding the solutions that exist for their concerns, such as embedded security and privacy guardrails, human-in-the-loop requirements, and strategies for re-allocating billable hours to other revenue-generating initiatives.

2 Misinformation

Early reports of AI hallucinations and unreliable outputs have left a lasting impression. These concerns, while valid in the early days of generative AI and particularly so in high stakes work such as law, accounting, and financial services, are often misapplied to today's more advanced and secure agentic systems. Professionals may not realize that today's strong agentic AI solutions operate in a controlled, enterprise-grade environment, far removed from the open web tools they're wary of. For example, agentic systems developed by Thomson Reuters are grounded in authoritative content curated for professional use. This reduces the risk of hallucinations and helps ensure that outputs are aligned with legal, regulatory, and industry standards.

3 Workflow integration anxiety

Even organizations that have begun experimenting with AI often struggle with embedding it into their workflows. The fear is that AI will disrupt established processes or require significant retraining. However, agentic AI is designed to work like a junior professional that is capable of efficiently handling multi-step work (but always under human oversight).

4 Lack of clear differentiation

With every vendor claiming to offer "agentic AI," organizations are understandably skeptical and may lack a framework against which to evaluate their options. What separates a true agentic system from a glorified chatbot comes down to orchestration, evaluation, and memory. A legitimate agentic system uses large language models (LLMs) to call tools in a loop, make decisions, and adapt over time—something most off-the-shelf solutions can't do.

5 Endless debate over implementation vs. ROI

There is a realization within companies that inaction has a cost, but initial adoption will stall because they can't determine which applications will generate the most return on investment (ROI) the fastest. They resist the idea of starting small, in low-risk environments, which is often the key to building confidence and forward movement within the company.

This last point is essential. Because employees at any organization have different roles and varying levels of familiarity with agentic and generative AI, the easiest inroad to adoption is starting with more accessible tasks. From there, professionals can utilize more sophisticated applications that generate even more benefits and ROI.

Accessible use cases for legal, tax, and risk professionals

Bridging the gap between wanting to implement agentic AI and making it a part of your everyday work begins with understanding what specific tasks the technology can assist with (determining use cases) and how that fits into standard workflows.

An essential underlying component of agentic AI, LLMs, is what makes these tools so accessible. LLMs use natural language, which means users can ask agentic AI questions or give directions (known as queries or prompts) in the same way they would when speaking to another person. Below are some examples of prompts a user might assign to their agent:

ⓘ Legitimate, professional-grade agentic AI can draw information from multiple, disparate, and trusted sources—for example, an organization's proprietary data, vetted industry-specific tools like Westlaw for legal research, CLEAR for investigative and compliance insights, official public databases, and reliable vertical news sources. A reliable agentic AI solution, such as CoCounsel, will automatically cite and link to the sources of information throughout its responses so users can quickly access the original information.

Reviewing a document

Scan, retrieve, analyze, summarize, and correct information

> **Legal:** Review this NDA and extract any clauses regarding obligations of the receiving party. I also want you to flag risky or missing terms, summarize consequences of breach, and suggest edits based on state precedent or company policy.
>
> **Tax:** Can you check for any potential errors in this document? Also, ensure correct deductions/credits, and analyze year-over-year changes. Make suggestions for improvement and explain why you made those suggestions.
>
> **Risk:** For this document of customer profiles, assess the clients most at risk of defaulting on their loans. Prioritize them by risk level and explain your reasoning.

Checking compliance

Interpret information within the context of various regulations, requirements, and policies

> **Legal:** Check this service agreement for compliance with California labor law, especially regarding non-compete and at-will employment clauses. Also, highlight any indemnity clauses that exceed our company's standard liability threshold.
>
> **Tax:** Review this fixed asset register for depreciation methods and confirm they comply with MACRS.
>
> **Risk:** Review this batch of wire transfers for violations of our bank's anti-money laundering thresholds. Is there any SAR-worthy activity worth flagging?

Researching and investigation

Find, analyze, and organize essential information

> **Legal:** What are the current FTC guidelines for advertising disclosures on social media?
>
> **Tax:** What are the tax implications of intercompany loans under U.S. and OECD transfer pricing rules?
>
> **Risk:** Cross-reference these IP addresses and phone numbers with known fraud cases or scam operations.

Drafting a document

Prepare customized and accurate documentation and templates

> **Legal:** Draft a terms of service agreement for a subscription-based mobile app, including a dispute resolution clause.
>
> **Tax:** Create a cover letter for submitting an amended tax return to the IRS with an explanation of the changes.
>
> **Risk:** Generate a fraud investigation report including findings, timelines, and recommended remediation steps.

Training and knowledge retrieval

Help junior associates advance their knowledge in particular areas and better understand processes faster

> **Legal:** Should I be using a unilateral or mutual NDA for this case? Explain the difference and why I should use one and not the other.
>
> **Tax:** How can I check whether a client is eligible for the R&D tax credit? Break the process down into steps.
>
> **Risk:** Walk me through how to perform a basic third-party risk assessment using a standard checklist.

Now that you're familiar with the types of tasks agentic AI can complete and how to initiate them, let's look at what entire processes might look like for legal, tax, and risk use cases, including the human-in-the-loop interaction.

Why are humans essential to agentic AI workflows?

Even though agentic AI can independently execute complex tasks concurrently, humans remain the drivers and safekeepers of the workflow, particularly in certain phases. This is referred to as human-in-the-loop interactions. Humans choose which prompts or queries to use, refine and iterate on those prompts by revising their requests or changing direction of the task, and review the final product to ensure accuracy, validity, and relevance.

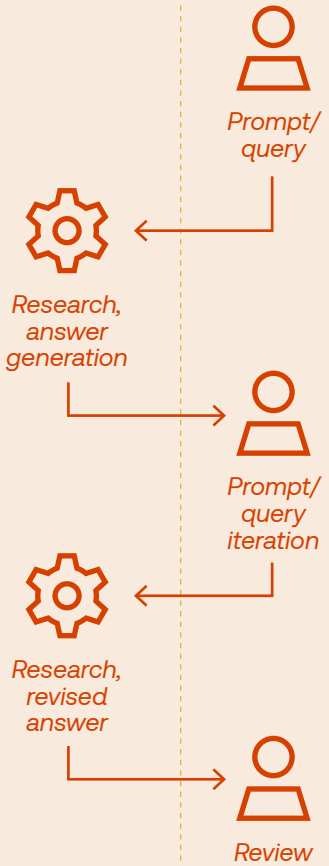

Legal use cases

Legal professionals are under increasing pressure to deliver faster, more accurate work while navigating growing complexity across jurisdictions, practice areas, and client demands. Agentic AI offers a transformative solution by automating multi-step legal workflows while keeping attorneys in control. Unlike traditional automation or standalone generative tools, solutions powered by agentic AI such as CoCounsel Legal can plan, adapt, and execute tasks across platforms, surfacing insights that might otherwise be missed. This enables law firms to scale expertise and deliver higher-value outcomes with greater speed and confidence.

Use case #1: Policy drafting

One of the most tangible examples of how agentic AI can transform legal workflows is policy drafting. Today, when an attorney needs to draft an employee policy, such as a data privacy or workplace conduct policy, they must manually locate a relevant template, research applicable state laws, and tailor the language to their organization's specific needs. This process can take hours or even days, especially when compliance across multiple jurisdictions is involved.

With agentic AI, this workflow becomes dramatically more efficient. The attorney begins by entering a natural language prompt into the system, such as, "Draft a workplace marijuana policy for a company with 500 employees operating in California and Colorado." The AI agent then initiates a multi-step process:

1. Planning and research	The agent identifies the relevant legal frameworks in both states, pulling from Practical Law's jurisdiction-specific guidance and Westlaw's primary law content. It determines which provisions are required, optional, or prohibited in each jurisdiction.
2. Template selection and customization	Based on the user's input and legal research, the agent selects a base template and begins customizing it. It incorporates the user's parameters (e.g., company size, location, policy goals) and adjusts the language accordingly.
3. Draft generation	The agent produces a first draft of the policy, complete with citations to the underlying Practical Law content that informed each clause. This "show your work" approach allows the attorney to verify the legal basis for each provision.
4. Human-in-the-loop iterative refinement	If the agent identifies gaps, such as missing information (e.g., "Specify whether the company allows off-duty use"), it prompts the user for clarification. This ensures the final output is tailored, accurate, and aligned with the firm's needs.

5. Human-in-the-loop review	Before finalizing the draft, the agent presents a summary of its plan and key decisions to the user. The attorney can review the draft, request revisions (e.g., “Add a section for unionized employees”), or provide missing information.
6. Training and knowledge transfer	For junior attorneys, this workflow doubles as a learning experience. The citations and linked content provide context and rationale, helping them understand not just what the policy says, but why it says it.

This same agentic framework can be extended to other legal workflows, such as preparing litigation documents. In each case, the agent handles the heavy lifting—research, synthesis, drafting—while the attorney guides, reviews, and validates the output.

Use case #2: Incorporating a company

Incorporating a company is a multi-step legal process that involves selecting the appropriate entity type, preparing and filing formation documents, drafting bylaws, registering for tax IDs, and ensuring compliance with jurisdiction-specific regulations. Traditionally, this process requires attorneys to manually coordinate across multiple systems, reference various legal sources, and manage client communications.

With agentic AI, the process becomes orchestrated, adaptive, and significantly more efficient. The workflow might look like this:

1. Initial prompt and planning	The attorney begins by entering a prompt such as, “Incorporate a Delaware C-Corp for a SaaS company with three founders and remote employees in five states.” The agentic AI system interprets the request and generates a step-by-step plan, outlining the necessary filings, documents, and compliance checks.
2. Tool orchestration and content retrieval	The agent accesses internal tools and trusted content sources such as Practical Law and Westlaw to retrieve jurisdiction-specific requirements, entity type comparisons, and filing procedures. It selects the appropriate templates and legal forms based on the user’s input.
3. Human-in-the-loop review of the plan	Before proceeding, the agent presents the incorporation plan to the attorney, including the rationale for each step and links to the legal sources used. The attorney can approve the plan, request modifications, or provide additional context (e.g., “Include a clause for founder vesting schedules”).

4. Execution with adaptive prompts	As the agent begins drafting documents and preparing filings, it may encounter missing information, such as the company's registered agent or the founders' equity split. It pauses and prompts the attorney to supply the missing details, ensuring accuracy and compliance.
5. Draft generation and validation	The agent generates the incorporation documents, including the certificate of incorporation, bylaws, and initial board resolutions. Each document includes citations to the legal sources that informed its content, allowing the attorney to "trust but verify."
6. Iterative refinement and filing	The attorney reviews the drafts, makes any necessary edits, and instructs the agent to proceed with filing. The agent can then prepare the documents for submission, depending on integration capabilities.
7. Training and knowledge transfer	For junior attorneys, this workflow serves as a training tool. By reviewing the agent's plan, citations, and drafting logic, they gain insight into the incorporation process and the legal reasoning behind each step.
8. Future expansion	As agentic capabilities mature, this workflow could expand to include post-incorporation tasks such as issuing stock, setting up cap tables, and registering for state-specific business licenses, further reducing manual effort and increasing consistency.

This scenario highlights how agentic AI can turn a traditionally fragmented and manual legal process into a streamlined, intelligent workflow that takes a fraction of the time. It also reinforces the importance of human oversight, not just for validation, but for guiding the AI with strategic decisions and nuanced judgment.

Tax and accounting use cases

One of the most promising applications of agentic AI in tax and accounting is in tax advisory services. This workflow, which traditionally involves a mix of manual research, spreadsheet modeling, and client communication, evolves into a guided, intelligent process that accelerates outcomes while maintaining professional discretion.

Use case #1: Tax strategy planning and execution

A CPA advising a small business client on how to reduce their tax liability for the upcoming fiscal year typically begins by reviewing the client's financials, identifying applicable tax strategies, modeling potential outcomes, and preparing a proposal. This process is time-consuming, error-prone, and heavily reliant on the advisor's ability to synthesize complex tax code and firm-specific data. With agentic AI, the workflow becomes structured, adaptive, and significantly more efficient:

1. Client intake and document review	The CPA uploads the client's financial documents, such as prior tax returns, income statements, and ownership structure, into the system. The AI agent reviews the documents, extracts relevant data (e.g., revenue, entity type, number of employees), and flags any missing or inconsistent information. The CPA then uploads that additional information, links to it, or responds directly to the agent with the needed data.
2. Strategy identification and eligibility assessment	The agent then evaluates the client's profile against a library of tax strategies (e.g., income shifting, entity restructuring, retirement plan contributions). It determines which strategies are applicable based on eligibility criteria and regulatory thresholds, using embedded logic derived from the agentic platform's content, such as Checkpoint.
3. Impact modeling and proposal drafting	For each viable strategy, the agent calculates the estimated tax impact and generates a draft proposal. This includes a summary of the strategy, projected savings, and a step-by-step implementation plan. The proposal is linked to authoritative sources, allowing the CPA to confirm the recommendations.
4. Human-in-the-loop review and customization	Before sharing the proposal with the client, the CPA reviews the agent's output. If needed, they can adjust assumptions, add client-specific context, or request clarification from the agent (e.g., "What assumptions were used in the income shifting model?"). The agent responds with citations and reasoning, enabling the CPA to better understand and explain rationale to the client.

5. Guided execution workflow	Once the client approves a strategy, the agent initiates a guided workflow to implement it. For example, if the strategy involves forming a new LLC, the agent outlines the required steps (filing articles of organization, updating payroll systems, notifying stakeholders, etc.) and tracks progress. At each step, the agent prompts the CPA for inputs or confirmations (e.g., "Upload the signed operating agreement").
6. Compliance and documentation	Throughout the process, the agent ensures that all actions are documented and compliant with relevant tax codes. It generates audit-ready records and flags any deviations from standard procedures.
7. Training and knowledge transfer	For junior accountants, this workflow serves as a built-in training tool. The agent's explanations, citations, and structured guidance help them understand not just what to do, but why—accelerating the learning curve and reducing reliance on senior staff.

Here, we see how agentic AI can turn a complicated advisory process into a task that requires far less effort. By pairing automation with expert content and human oversight, tax professionals can deliver faster, more accurate, and more strategic advice while maintaining the trust and rigor their clients expect.

Use case #2: Document review and compliance checks in tax advisory

In addition to planning and executing tax strategies, CPAs often need to review a client's financial documents to ensure completeness, consistency, and compliance with applicable tax regulations. This process is largely manual and includes comparing multiple documents, cross-referencing with tax codes, and identifying missing or conflicting information. With agentic AI, this document review process becomes more intelligent, proactive, and guided:

1. Document upload and preprocessing	The CPA uploads a set of client documents, such as W-2s, 1099s, prior returns, and financial statements, into the system. The AI agent scans and extracts relevant data points, such as income, deductions, and entity structure.
2. Automated review and flagging	The agent reviews the documents for completeness and consistency. It flags missing forms (e.g., "No Schedule C found for reported self-employment income") or inconsistencies (e.g., "Income reported on 1099 does not match bank statement deposits").

3. Compliance validation	The agent checks the extracted data against current tax regulations using embedded logic and authoritative content from Checkpoint. It identifies potential compliance issues or opportunities (e.g., “Client may qualify for R&D tax credit based on expense patterns”).
4. Human-in-the-loop interaction	When the agent encounters ambiguous or incomplete information, it prompts the CPA for clarification (e.g., “Please confirm whether the client has any dependents listed on their return”). The CPA can upload additional documents or enter responses directly.
5. Summary report generation	Once the review is complete, the agent generates a summary report outlining findings, flagged issues, and recommended next steps. Each item is linked to the relevant source document and regulation, enabling the CPA to sign off the output or return to the agent for revision.
6. Training and knowledge transfer	The agent’s explanations and citations help new employees understand the rationale behind each compliance check and how to interpret tax documentation more effectively.

This scenario demonstrates how agentic AI can enhance accuracy, reduce manual effort, and improve audit readiness in tax and accounting workflows. By combining document intelligence with regulatory awareness and human oversight, it enables professionals to deliver higher-quality outcomes with greater confidence.

Risk and fraud use cases

Agentic AI can dramatically accelerate and enhance the risk and fraud investigation process by automating the data gathering, synthesis, and risk flagging steps, while keeping the investigator in control of key decisions.

Use case #1: Investigative due diligence for a high-risk entity at a bank

Risk and fraud professionals—in both corporations and government agencies—often face the challenge of conducting deep due diligence on individuals or entities that may pose reputational, financial, or legal risk. This process involves gathering data from disparate sources, verifying identities, uncovering hidden relationships, and finding red flags such as criminal records, sanctions, or financial irregularities. With agentic AI, this investigative workflow becomes more intelligent, iterative, and auditable:

1. Initial prompt and objective setting	The investigator begins by entering a natural language prompt such as: "Conduct a due diligence investigation on John Doe, a potential vendor based in Miami, with suspected ties to offshore entities." The agent interprets the objective and generates a multi-step plan to gather, analyze, and synthesize relevant data.
2. Data aggregation and identity resolution	The agent accesses a wide range of public and proprietary data sources (court records, corporate registries, sanctions lists, and news archives) to build a comprehensive profile. It uses entity resolution techniques to distinguish between individuals with similar names and confirm identity through cross-referenced attributes like date of birth, address history, and known associates.
3. Relationship mapping and risk flagging	The agent constructs a network map of the subject's known business affiliations, associates, and historical transactions. It flags potential risks such as links to sanctioned entities, prior fraud investigations, or shell companies in high-risk jurisdictions. These insights are surfaced in a visual dashboard or structured report.
4. Human-in-the-loop review and refinement	At key decision points, the agent prompts the investigator to validate assumptions or provide additional context. For example, if the agent finds multiple John Does in Florida, it may ask, "Is your subject associated with XYZ Holdings LLC?" The investigator can confirm or correct the path, ensuring the AI stays on track. If the investigator doesn't know, the agent can suggest alternative paths to uncover further information.

5. Narrative report generation with source attribution	Once the investigation is complete, the agent generates a narrative report summarizing findings, supported by citations and links to original source documents. This allows the investigator to verify each claim and ensure transparency in how conclusions were reached.
6. Iterative exploration and scenario testing	If new information emerges, such as a tip about a related entity in Panama, the investigator can prompt the agent to expand the scope. The agent adapts its plan, retrieves new data, and updates the report accordingly, maintaining a clear audit trail of changes.
7. Training and knowledge transfer	For junior analysts or new team members, the agent's explanations, source links, and structured logic help them understand investigative best practices and regulatory frameworks. This helps them quickly become familiar with ways to overcome roadblocks or useful ways to find additional information.

The above scenario demonstrates how agentic AI can enhance the speed, depth, and reliability of risk investigations. By automating the labor-intensive aspects of data gathering and synthesis—all while keeping the human expert in control—agentic systems empower professionals to make faster, more informed decisions with confidence.

Use case #2: Investigative lead development for law enforcement

Police investigators are often tasked with identifying and assessing individuals or entities that may be involved in criminal activity, fraud, or public safety threats. These investigations usually involve pulling together fragmented data from multiple sources, such as arrest records, court filings, social media, business affiliations, and geographic movement patterns, to build a comprehensive picture of a subject's background and potential risk.

1. Initial prompt and objective setting	An officer begins by entering a prompt such as: "Investigate Marcus Taylor, a suspected gang affiliate operating in Atlanta and Dallas, with possible ties to synthetic drug distribution." The agent interprets the objective and generates a multi-step plan to gather relevant data, assess risk, and surface actionable leads.
2. Data aggregation and identity resolution	The agent pulls from a wide array of public and proprietary sources including arrest records, incarceration history, known aliases, vehicle registrations, and property ownership. It uses entity resolution to distinguish between individuals with similar names and confirm identity through cross-referenced attributes such as date of birth, known associates, and prior addresses.

3. Network mapping and behavioral pattern analysis	The agent constructs a network of known associates, flagged addresses, and prior incidents. It identifies patterns such as frequent travel between high-risk areas, connections to previously investigated individuals, or overlapping phone numbers and vehicles. These insights are visualized in a relationship map or structured report.
4. Human-in-the-loop review and decision points	At key junctures, the agent prompts the officer for clarification or confirmation. For example, if multiple individuals named Marcus Taylor are found, the agent may ask, "Is your subject associated with a 2021 arrest in Fulton County?" The officer's input helps refine the investigation and ensure accuracy. If the officer doesn't know, they can ask the agent for more details about that arrest to make a determination.
5. Narrative report generation with source attribution	Once the investigation is complete, the agent generates a narrative report summarizing findings, including citations and links to original source documents. This allows the officer to verify each claim and remain clear about how conclusions were reached.
6. Scenario expansion and real-time monitoring	If new intelligence emerges, such as a tip about a recent arrest or a new alias, the officer can prompt the agent to update the narrative and any ensuing investigative leads to keep them as up to date as possible. The agent adapts its plan, retrieves new data, and updates the report accordingly. In future iterations, the agent could also monitor and alert for real-time updates, such as new bookings or court filings.
7. Training and knowledge transfer	For newer crime analysts, the agent's explanations, source links, and structured logic help them understand investigative best practices, legal thresholds, and how to interpret complex data sets.

This real-life use case demonstrates how an AI agent can support law enforcement by reducing the time and effort required to build a case file while improving the depth and reliability of the analysis. By automating the labor-intensive aspects of data gathering and synthesis—and prompting for human judgment at critical points—agentic systems can help law enforcement keep up with increased incidents and more sophisticated criminals.

Strategies for success: Implementing AI with purpose

Agentic AI is a force multiplier, but success doesn't come from adopting the technology alone. Organizations must implement it with intention, align it with company strategy, and measure its impact in ways that go beyond surface-level efficiency. To do all this, companies should use several tactics that will improve success in the short- and long-term.

First, focus on **starting small** and thinking big. Begin with familiar, high-impact use cases—summarizing depositions, analyzing tax data, automating compliance checks—and use those early wins to build confidence and momentum. Don't view these steps as pilot projects: consider them strategic footholds. The key is to tie adoption directly to business priorities and roll out agentic workflows with a clear plan for change management, training, and support.

This kind of transformation doesn't happen in silos. The most successful implementations are **cross-functional**. They involve leadership, IT, pricing, and talent management working together to define what success looks like and how to get there. Those that treat agentic AI as an organization-wide initiative and not solely as a tech upgrade are the ones seeing the greatest returns. In fact, legal professionals expect to free up nearly 240 hours per year through AI adoption, unlocking an average annual value of $19,000 per professional, according to the 2025 Future of Professionals Report. Across the U.S. legal and tax sectors, that translates to a $32 billion opportunity.

22%
of organizations have a tailored strategy in place*

43%
are adopting AI without a strategy*

54%
of employees feel they have sufficient input into how organizations plan to use AI*

When thinking about ROI, it's natural to default to first thinking about time saved. Take it further and consider how your strategy will **use returned time**. Some professionals are using it to take on more work, others to shift toward fixed-fee billing models, and still others to deepen client relationships. The most forward-thinking organizations are measuring success through a combination of leading and lagging indicators: time saved per task, increased capacity, usage metrics, and ultimately, financial outcomes such as improved margins, increased revenue per professional, or greater client value realization. The companies that have found the most success with agentic AI both track these metrics and have a strategy for turning efficiency into value.

Choosing **the right partner** is just as critical as choosing the right use case. Organizations need more than a vendor; they need a collaborator who brings domain expertise, robust security, and a long-term commitment to innovation. Don't think only in terms of the agent but about the entire ecosystem surrounding it (content, tools, and guidance).

* Future of Professionals Report, Thomson Reuters, 2025.

For this reason, teams are increasingly looking to providers who can offer the training, support, and integration needed to make this technology work in the real world. The pace of change is fast, and the cost of inaction is rising. Thirty percent of professionals believe their firms are moving too slowly on AI, even as 80% expect it to have a transformational impact within five years. The firms that will thrive are those that embrace adaptability, foster a culture of experimentation, and build systems that are designed to evolve. The goal isn't to resist change—it's to design for it.

The AI Success Pyramid

Let's examine how each of these four layers can help organizations achieve the maximum benefits from their AI use.

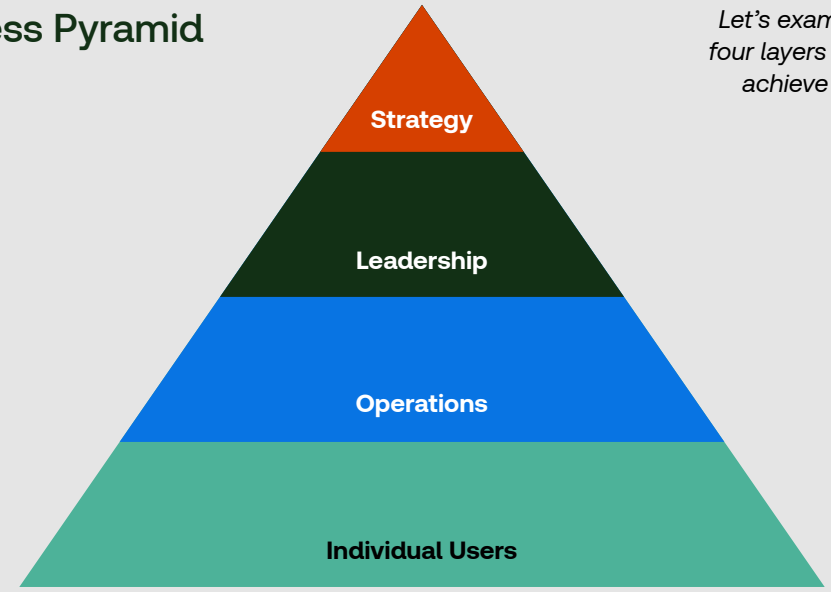

The Thomson Reuters 2025 Future of Professionals Report outlines a four-layer pyramid that defines the organizational foundation for successful AI adoption. At the top is strategy, where organizations articulate a clear vision for how AI will drive value across the business. This strategic clarity must be championed by leadership, the second layer, which plays a critical role in setting priorities, allocating resources, and modeling adoption. The third layer, operations, ensures that the right infrastructure, workflows, and governance are in place to support scalable implementation. At the base of the pyramid are individual users—the professionals whose daily engagement with AI tools ultimately determines the success of the initiative. Each layer reinforces the others, creating a cohesive framework that enables teams to move from experimentation to enterprise-wide transformation.

Agentic AI is not a silver bullet, but it is a powerful tool. When implemented thoughtfully, it can transform workflows, elevate client service, and unlock new levels of performance. Those that succeed won't be the ones that simply adopted it the fastest, they'll be the ones that also adopted it in the smartest way.

Reinvent your workday with a trusted AI agent

Thomson Reuters offers the most comprehensive and future-ready AI solution available for professionals who are looking to evolve under the weight of growing complexity and role expectations. It empowers professionals to work more accurately, thoroughly, and quickly by orchestrating entire workflows across research, drafting, analysis, and review.

Thomson Reuters isn't just ready for what's next—it's building it. That's why professionals who begin using CoCounsel and CLEAR today will be ready to apply agentic AI to more advanced use cases as new features and capabilities emerge.

CoCounsel Legal is the most advanced single AI tool for quickly and accurately completing complex legal work—and is the AI solution of choice for 78% of the Am Law 100 and the entire U.S. federal court system.

- Applies guided agentic AI across the full legal workflow
- Integrated with trusted Westlaw and Practical Law content
- Developed by 1,000+ AI, machine learning, and data science experts
- Backed by 225+ years of legal expertise
- Only AI legal assistant with ISO 42001 certification
- Includes citations for full transparency and review
- Continuously evolving with partners such as Noetica and DeepJudge

Learn more →

CoCounsel for tax, audit, and accounting professionals transforms core workflows with intelligent support for research, document analysis, advisory work, and audit testing.

- Combines trustworthy answers, automation, and firm knowledge seamlessly
- Built on trusted IRS code, Checkpoint, and firm sources
- Includes plain-language explanations with citations for full transparency
- Integrates seamlessly with Checkpoint and Excel workflows
- Fully cloud-native with secure controls for remote collaboration
- Automates real work from memos to client file reviews
- Eliminates manual re-entry, formatting, and document chasing
- Connects policies, files, and guidance in searchable workspace

Learn more →

CLEAR empowers risk and fraud professionals to make data-driven decisions and conduct faster, more accurate investigations by aggregating and analyzing data from thousands of public and proprietary sources.

- Connects to thousands of public and proprietary data sources
- Uses AI to resolve identities and uncover hidden relationships
- Flags potential threats with real-time alerts and risk indicators
- Visualizes networks of people, businesses, and assets
- Provides audit-ready transparency with source-linked results
- Speeds up investigations while maintaining accuracy and control
- Trusted by law enforcement, corporations, and public agencies

Learn more →