KPMG

Seize the future

The agentic shift in SOX compliance

The future of SOX is here—and it's powered by agents

Attention SOX practitioners: the landscape is shifting beneath our feet. The traditional ways of managing Sarbanes-Oxley (SOX)—with manual processes, endless spreadsheets, and constant back-and-forth with control owners—are no longer sufficient in the face of growing complexity and mounting pressure.

But there's good news: a new breed of technology is emerging that holds the key to not only meeting these challenges, but also unlocking a new era of efficiency, effectiveness, and strategic value for SOX programs. This technology is agentic artificial intelligence (AI), and it's poised to fundamentally transform the way we approach SOX compliance.

The potential applications for agents in SOX are vast. From automatically collecting and organizing evidence, to continuously monitoring controls, to surfacing insights and anomalies—agents can take on the most tedious, time-consuming aspects of SOX and perform them faster, more consistently, and more thoroughly. But the impact of agents goes beyond just efficiency gains. By enhancing the quality, reliability, and breadth of SOX testing, agents can fundamentally elevate the level of assurance you provide. They can help you proactively identify and mitigate risks, respond more nimbly to changes in the business, and provide deeper, more actionable insights to stakeholders. In short, agents can help you transform SOX from a necessary compliance activity to a strategic advantage.

In the following pages, we'll take a closer look at what exactly agents are, explore some key use cases for SOX, and provide a practical roadmap for getting started with agents in your own program. We'll also examine how agents fit into the broader landscape of SOX technology and methodology while discussing the implications for the future of the SOX profession.

The age of agents is dawning, and with it comes a once-in-a-generation opportunity to reimagine SOX compliance. The question is: Will you be at the forefront of this transformation or watching from the sidelines?

Read on to learn how you can harness the power of agents to take your SOX program to the next level.

What is an agent?

The now ubiquitous generative AI (GenAI) solutions can provide answers based on a user prompt. One of the things that makes these tools so impressive is their ability to take the plain language of the user and create humanlike responses that reflect contextual understanding and reasoning capabilities. Yet, even with this impressive potential, early GenAI tools are still just "answer machines." They cannot take action on your behalf. Agents are different from traditional GenAI in that they have the ability to perform a wide variety of tasks rather than simply generate content.

Consider this simple example. A controls tester wishes to write a detailed email to a control owner asking for information about some evidence. With a traditional GenAI tool, they would type their prompt with the relevant details, and the GenAI tool would produce a draft of the email. The controls tester would typically then have to open their email client, create a new email template, copy and paste the text of the email, look up the person's email in the company's global address book, review and revise the email, and hit send. An agent could take the plain language of the prompt and perform all of these steps on behalf of the user—or if the user prefers, the agent could perform all of the steps except pressing the send button so as to allow for a human-in-the-loop review.

The potential for agents to transform SOX compliance cannot be overstated. Imagine a future where the most time-consuming, repetitive tasks in your SOX program are fully automated, freeing up more of your team's time to focus on critical activities like risk assessment, control design, and root cause analysis. Picture a world where agents continuously monitor your controls, proactively identifying issues and suggesting remediation steps before they become deficiencies. Envision a SOX program that's not just more efficient, but more effective—one that provides real-time assurance, actionable insights, and strategic value to the business. This is the promise of agents—and it's not some far-off vision, but a reality that's already starting to take shape today. As agents become more sophisticated and integrated into our SOX workflows, they will fundamentally change the way we work and the value we deliver. They will elevate the role of the SOX practitioner, and they will help us build SOX programs that are not just fit for purpose, but fit for the future. The question is not if agents will transform SOX, but when. And each SOX leader must decide to lead the charge now or play catch-up later.

Agents are driving productivity today

Every major cloud provider, productivity software suite, and enterprise resource planning system offers agentic capabilities today. There are also countless start-ups and mid-sized technology providers that are offering several types of agentic solutions. The adoption of agents to perform common SOX program tasks is not yet widespread, but the technology is widely available.

The biggest challenge with many attempts to automate SOX tasks in the past was that they failed to provide a return on investment (ROI). That is to say, building, testing, and deploying automated control testing solutions using legacy technologies might require 50 or more hours per control. Yet, many of these same control tests could be performed manually in about 10 hours.

The ROI was further diluted by the fact that even minor changes in control performance or evidence formats would often prevent these attempts at automation from running successfully in later periods.

AI-enabled agents can overcome many of these barriers to achieving ROI. The reasoning capabilities of GenAI models can help to adapt to the complexities of changes in documentation, data, and other evidence. Furthermore, the most common orchestration tools have robust integrations with many applications that allow for a faster, more seamless development and deployment of agents. And these capabilities are improving rapidly. Now is the time to begin your agentic transformation.

The TACO framework for agentic AI

As we consider how to use agents to improve SOX program outcomes, it is useful to think about the types of activities that agents are well suited to. For this, we use the TACO framework, which contemplates four primary agent types.

Taskers focus on accomplishing singular goals by breaking them into structured, repeatable tasks, making them easy to deploy, scale, and monitor. Examples for SOX practitioners include:

- Using a vision model to "read" a document and compare the contents of the document to predefined business rules.
- Extracting key terms from a large number of contracts and writing those terms into a table for use in data analysis.
- Documenting the contents of the data set and reviewing for data quality issues.

Automators handle goals that require integration across multiple enterprise applications, automating the performance or testing of end-to-end business processes and

cross-functional workflows. Examples for SOX practitioners include:

- Extracting role and user data from multiple systems, preparing data for comparison, and performing cross-system segregation of duties analysis.
- Comparing spending delegation of authority grants in the purchasing system(s) to enterprise policies and identifying anomalies for follow-up.

Collaborators act as AI teammates, working contextually and closely with human operators, enhancing productivity by learning from interactions and refining recommendations. Examples for SOX practitioners include:

- Supporting the evaluation of management review controls based on similar controls, related exceptions, and scanning control evidence for possible red flags to provide suggestions to the tester about questions to ask.
- Extracting and providing evidence sources based on similar or historical controls, systems assessed, and contextual metadata (e.g., fiscal period, control frequency) then renaming files based on preferred naming conventions and storing in predetermined file repository.

Orchestrators involve *multiagent ecosystems* where agents collaborate not only with humans but also coordinate with other AI agents to achieve complex tasks at scale while dynamically adapting to real-time changes. Examples for SOX practitioners include:

- Scanning servers and comparing actual configuration to required benchmark configuration, logging the discrepancy, and self-healing the exception by resetting the server configuration to meet requirements.
- Sending reminders to control owners with outstanding remediation plans and, based on the response received, the agent performs preliminary review of evidence of the remediation to determine the likelihood that the issue has been satisfactorily resolved.

Getting started with agents in your SOX program

When beginning your journey with agents, indeed with any technology enablement program, it is critical to establish clear objectives. The most successful programs have objectives that are measurable and aligned on outcomes, rather than activities. For instance, "Reduce control testing hours by 10 percent" is a much more useful goal than "Use agents on 20 controls." There are many types of outcome-oriented goals. You may choose to focus on hours or cost reduction, timeliness and repeatability, insight and process improvement, lessening control owner burden, improved assurance or coverage, and more. This goal should make clear which agentic use cases are in line with your objectives and which ones are not. This will be a critical element of your prioritization.

While not essential, it is advisable to thoroughly review your SOX-relevant business processes, in-scope systems, and controls catalog and look for opportunities to rationalize controls. For obvious reasons, there is little value in automating tasks that are duplicative, lower value, or otherwise unnecessary. While this can be a nontrivial effort, it is a wise starting point to optimize your investment and may bring you closer to your goal before you bring any agentic capabilities to bear.

Agents in the First Line of Defense

One of the biggest challenges to automating a SOX program is that many financial reporting controls have manually performed components. Control performance often takes significantly more time than control testing. While it is certainly possible for SOX practitioners to realize cost reductions using agents to test such manual controls, the manual performance of SOX controls may be an even better target for agentic automation approaches. Controls can be more preventive, less prone to failure, and more cost effective when they are technology-enabled, i.e., using workflow, data integration, system configuration, etc. Moreover, it is generally much easier to test automated controls, which further reduces the cost of SOX compliance and boosts the ROI on automating control performance. Financial reporting teams are well-advised to focus on finding ways to embed agents and other appropriate technologies into their processes.

Determine which agents to build and in what order

Your prioritization should also consider the degree of difficulty associated with the task you wish to build an agent to perform. These considerations are similar to those for automation and include the quantity and variety of data/evidence sources, the complexity of task, the degree of manual intervention on the performance of the underlying control/process, and several other factors.

At the outset of your program, it is also essential to consider the AI and agentic capabilities of your organization. Work with your AI governance team and information technology (IT) to understand the process to develop agents. It will help the discussion if you bring a few clearly defined examples of agents you wish to build. Armed with this information, you should be well positioned to determine where it makes sense to deploy agents and how to prioritize which agents to build. Of course, building the agents is simply the first of many steps in the new process.

Planning the deployment of your agents

Beginning to introduce agents—or any transformational technology—to your SOX program will likely raise methodological questions. These can be fairly trivial, such as how you document the work of the agent. For instance, does the agent need to apply tick marks to work papers? But, the methodological implications can also be more significant. Does the ability to test controls more frequently change your overall testing schedule? Does the ability to test a full population change your criteria for concluding on the effectiveness of the control? How does the use of agents impact your external auditor reliance strategy? The answers to these questions and many more will depend on the nature of your program and your goals. And while not all of these questions require answers during the early stages of experimentation, they should not be unduly delayed as they have implications for both how the agents are built and how your program evaluates the agents' output.

Establishing expectations with team members and other stakeholders is also an important step in the successful deployment of agents. They need to clearly understand the objectives of the new processes and how they will be impacted. This may mean quelling concerns about their future responsibilities, providing guidance on how to follow up on the outputs of agentic tests, or communicating details about new human-in-the-loop review procedures. After you have made decisions about the objectives of the program, you can then create a change strategy that is commensurate with the magnitude of behavioral change expected by your agentic program. As you roll out the agents, communicate and celebrate the successful outcomes of the change and the stakeholder behaviors that facilitated the success. This can help to cement the right behaviors and excite your team to even more progress.

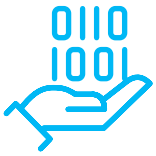

Building your first SOX program agents

The SOX agents that will yield the most value to your organization will be dependent on your strategy, objectives, and capabilities. The examples below all come with numerous assumptions about technical considerations, such as system and data availability, network architecture, etc., as well as process considerations, such as the extent to which processes are manual and consistent. Nevertheless, the following agentic use cases have broad applicability and merit strong consideration by most SOX teams.

Evidence collection

Collecting the evidence required to test SOX controls is time consuming for both control owners and testers. It creates bottlenecks that can add duration to tasks. And sometimes, control owners provide incorrect or incomplete information. Enterprise class agentic capabilities allow agents to access shared storage areas and applications to retrieve documents, data, and other common forms of evidence. Agents could be orchestrated to perform any or all of these tasks:

- Obtain populations from source systems
- Follow approved sampling procedures to select a sample
- Collect the evidence from the relevant source locations
- Save collected evidence in a predefined location
- Change the format or naming of evidence
- Determine the completeness of the information, including the presence of all required attributes
- Extract certain features from evidence and save in predetermined, e.g., tabular format.

Such agents will be efficiently built and configured on a control-by-control basis as the SOX team tests each control. However, it is also possible to create templated agentic workflows that allow control owners to configure their own agents to source their evidence. This may be useful when the control owner wishes to configure the agent at the time the evidence is being created.

Walk-through performance and associated documentation

Performing process walk-throughs can be a complex process. GenAI has been useful for the last two years to do things like help with taking existing inputs like call transcripts and notes, standard operating procedures, and system documentation and converting this to process narratives. Agentic capabilities can further expand this to facilitate the process of performing and documenting walk-throughs from start to finish, including:

- Scheduling meetings with all relevant stakeholders
- Drafting detailed agendas and background information to prepare control owners and others for the nature of the meeting
- Compiling all of the inputs required to prepare for a walk-through into a comprehensive interview guide
- Providing detailed briefings to the SOX team in advance of the meeting in their preferred learning style, e.g., a short podcast, an infographic, a memo, a spoken dialogue with the AI, etc.
- Transcribing calls or recordings of in-person meetings
- Extracting key process steps from these recordings and other relevant information to create or update process narratives or flow charts, including annotation of risk and controls
- Proposing updates to risk and control matrices based on walk-through outputs
- Recommending updates to accounting manuals or other standard operating procedures based on process changes or variations identified during walk-throughs.

While some of these tasks can be performed with nonagentic GenAI capabilities, agentic capabilities allow these tasks to be more seamlessly orchestrated without prompt engineering and swivel chair activities at each step. Moreover, it is not difficult to imagine next steps in which the AI could simply perform the interview with the control owner and report back to the SOX team. This may seem like science fiction, but the reality is that current technology is so close to performing this task well that the limiting factor for adopting such AI-enabled interactions is more likely to be audit standards and social norms than technical feasibility.

Integrated third-party risk and SOC report review

Reliance on third parties can introduce control dependency risks that must be continuously monitored and documented. Agents enable scalable and intelligent management of SOC reports and third-party risk. Agentic capabilities can:

- Automatically request the reports from the relevant parties
- Extract relevant control descriptions, testing results, exceptions, and user entity complementary controls
- Map those controls with internal IT general controls and business process controls
- Flag gaps, such as carve-outs, subservice organizations, and timing mismatches, to recommend compensating controls or enhanced testing
- Actively follow up on the collection of bridge letters, monitoring report updates, and documenting ongoing risk exposure for human assessment
- Aggregating third-party risk information to identify holistic themes across service providers or recurring issues and trends at individual vendors.

One of the attributes that makes this use case valuable is that very little human intervention is required outside of reviewing the results. The triggering event for such an agent could be something as simple as uploading a file to a specified folder or forwarding it to an email address set up for that purpose. The presence of a new document initiates the agentic workflow which in turn notifies the SOX team upon completion of the analysis.

SOX calendar and planning

Planning the SOX program execution calendar can be a time-consuming exercise, filled with competing priorities, bottlenecks, and other scheduling challenges. Agents can help by taking into account these issues and proposing an optimal SOX schedule to reduce the risk of slippage, improve collaboration, and use resources more efficiently. For example, agents can:

- Review timelines and completion dates from prior SOX cycles to identify areas and activities that have historically caused delays
- Account for blackout periods and other scheduling issues, such as period-end close timelines, audit committee meetings, holidays, planned system upgrades or outages, etc.
- Load balance activity levels across team members and consultants to balance workload, cost, and schedule
- Provide detailed timelines that account for dependencies and propose dates for tasks, such as kick-off meetings, control owner reminders, walk-through dates, document requests, controls testing, etc.
- Use schedules from other assurance providers, i.e., internal audit, enterprise risk functions, and other second line of defense teams, to coordinate on shared activities to remove duplication
- Monitor progress, provide interactive status reporting, send personalized reminders, and recommend schedule changes based on actual progress.

Agents are expected to play an increasingly large role in calendar management generally and project management specifically. And many of the current SOX program management platforms have capabilities that can facilitate program planning. However, the integration of enterprise calendars and program information from other stakeholders make agentic solutions to SOX planning compelling.

Controls testing

The single largest set of tasks in a SOX program is the testing of controls. Agents can help SOX practitioners perform control testing, sometimes performing all of the required steps while other times performing partial tests. Agents can be orchestrated to deal with the numerous small, but complex tasks required to test a control, including:

- Extracting relevant information from a wide range of source materials, such as unstructured documents and tabular data
- Write relevant information to special-purpose tables, lead sheets, or other data repositories
- Analyze extracted data and compare to defined attributes or benchmarks
- Review computer code that may be used in IT controls or to source information
- Document the results of work and proposed conclusions on control effectiveness
- Update the status of a control test.

It is worth noting that agents and GenAI have been used to test and monitor controls countless times by many organizations. However, this remains one of the single most difficult agentic solutions to perform at scale given the current state of technology. There are a few key reasons for this. Chief among them is the sheer variety of different controls and control types as well as the variability in the ways the performance of those

controls are evidenced. It is technically possible to build agents that will test nearly any control, but the volume and variety means that many SOX programs have significant effort before them. The second matter is that many controls require the application of complex logical processes that are often not clearly defined. Determining whether a management review control was performed effectively, for instance, rests on the tester's ability to assess the quality of the review. Distinguishing between a thoughtful review and perfunctory one is a task that can push boundaries for agents.

Despite these challenges, SOX practitioners needn't use agents to test their entire controls catalog. The technology will evolve as will the ability to use it. Growing enterprise investments in agentic AI make it increasingly imperative for SOX programs to begin the process of building agentic solutions for controls testing. SOX program leaders who take a wait-and-see approach may find themselves disrupted as technology advances beyond their strategic horizon.

Agents will redefine the future of SOX—are you ready for it?

The advent of agentic AI represents a watershed moment for SOX programs. The SOX leaders of tomorrow will not be those with the largest teams or the most resources, but rather those who most effectively harness the transformative power of agents. Agents are not just another tool in the toolbox—they are a fundamentally different way of working that will redefine what is possible in SOX compliance.

The business case for agents is undeniable. The efficiency gains, the enhanced quality and assurance, the cost savings, the ability to do more with less—the benefits are simply too compelling to ignore. And as agents become ubiquitous across business operations, SOX programs that fail to adapt will rapidly fall behind and struggle to keep pace.

Agents offer more than just an efficiency boost; they represent a significant opportunity to enhance the role of SOX. They provide the chance to become a more strategic partner to the business. They also offer the opportunity to inspire your employees by automating mundane tasks, allowing them to focus on more engaging work. Furthermore, agents present the opportunity to lead by example and set the standard for world-class SOX compliance.

Swarming your controls

As we consider the future of agentic AI, one of the emerging concepts that may be relevant to SOX is an "agent swarm." The idea is that many agents work together to perform a task or solve a particular problem. They may specialize, with each agent focusing on one small aspect of a task or problem. Alternatively, they may each be directed to take a slightly different approach to completing the same task then "voting" on the final answer. Imagine you create a primary agent that has the ability to prompt an agent swarm with 1,000 "sub-agents" all focused on testing a certain complex control. Because of the complexity of the control, some may take approaches that fail or get the wrong answer, but the majority will not. Agent swarms may prove to be more resilient in the face of complex or ambiguous tasks. Such systems exist primarily in academic and research settings today. There are several technical barriers to overcome before agent swarms can reliably and efficiently be deployed in an enterprise setting, but they offer a glimpse into a possible future for internal controls testing.

The time for incremental change is over. The age of agents is here, and it demands bold action. SOX leaders must have the vision to reimagine their programs, the courage to challenge the status quo, and the resolve to drive transformation.

KPMG LLP stands ready to help you seize this opportunity. Our distinguished experience in SOX, deep knowledge of agentic technologies, and established track record of driving innovation distinctly position us to guide you on this journey.

But the question is: will you lead the charge or watch from the sidelines? Will you be the disruptor or the disrupted? The choice is yours, but the stakes have never been higher. The future of SOX is here—and it will be defined by those with the vision to embrace the power of agents. Let's get started.

How KPMG can help

KPMG offers a wide range of solutions to support your SOX program, including:

Agentic readiness and technology strategy: Our SOX and technology practitioners work with your team to identify opportunities to deploy agents and other enabling technology that can help you accomplish your objectives. A clear strategy and a roadmap with prioritized use cases that meet your budget and technological capabilities can help you avoid pitfalls and progress more quickly.

Agent development and technology implementation: KPMG is a clear leader in implementing agentic AI technologies. With strong alliances across nearly every major cloud provider and software vendor, we know how to get the most out of your technology environment. Our Agentic Enablement team will help you move fast, reduce risk, and improve the performance of your SOX program.

SOX Program HealthCheck: We can assess your program and strategy against peers and objective benchmarks to find opportunities to improve program cost, efficiency, and collaboration with the business.

SOX cosourcing: Our experienced practitioners are recognized leaders in SOX compliance. With our deep industry experience, efficient delivery model, and high-quality client experience, you and your stakeholders can increase trust, manage budget effectively, and stay focused on your core business.

At KPMG, we believe that agents represent the future of SOX compliance—a future where SOX programs are more efficient, effective, and valuable than ever before. But we also recognize that the path to this future is complex and filled with uncertainty. That's why we're committed to being your trusted guide and consultant every step of the way.

Whether you're just starting to explore the potential of agents or you're ready to embark on a full-scale transformation, KPMG has the experience, knowledge, and tools to help you succeed. Let's work together to reimagine your SOX program and unlock the full power of agentic technologies.

Contact us

Charles King
AI in Internal Audit Leader
KPMG LLP
E: crking@kpmg.com

Sue King
SOX Offering Leader
KPMG LLP
E: susanking@kpmg.com

The authors would like to thank the following reviewers for their invaluable contributions to this article: Bernal Cantillo-Aguero, Rachel Horne, Parth Jhaveri, Richard Knight, and Dane Paulson.

Some or all of the services described herein may not be permissible for KPMG audit clients and their affiliates or related entities.

Learn about us: **kpmg.com**