Grant Thornton

Agentic AI: Future of Internal Audit

Tuesday, June 16, 2026

Speakers

Matthew Cassidy

Partner,
Internal Audit & SOX
Grant Thornton Advisors LLC
matt.cassidy@us.gt.com

Daryl Ray

Exp Manager,
Internal Audit & SOX
Grant Thornton Advisors LLC
daryl.ray@us.gt.com

Learning Objectives

1. Explain how agentic AI shifts the audit focus from outputs to system behavior and execution.

2. Identify where control exists in AI-enabled systems, particularly within the execution layer (permissions, approvals, and logging).

3. Recognize how audit approaches must adapt based on levels of AI autonomy and the associated business impact (risk) of AI-enabled activities.

4. Redefine Internal Audit's role in AI-driven environments by translating the concept of a "trusted advisor" into earlier engagement, behavior-level assurance, and delivering actionable insight where the business lacks visibility.

5. Define the future role of Internal Audit in AI-driven environments by distinguishing between "audit with AI" and "audit of AI," and explaining the shift from output validation to control architecture and system-behavior assurance.

Grant Thornton

Grant Thornton

Explain how agentic AI shifts the audit focus from outputs to system behavior and execution.

Grant Thornton

We are at the tipping point between experimentation and enterprise AI embedding

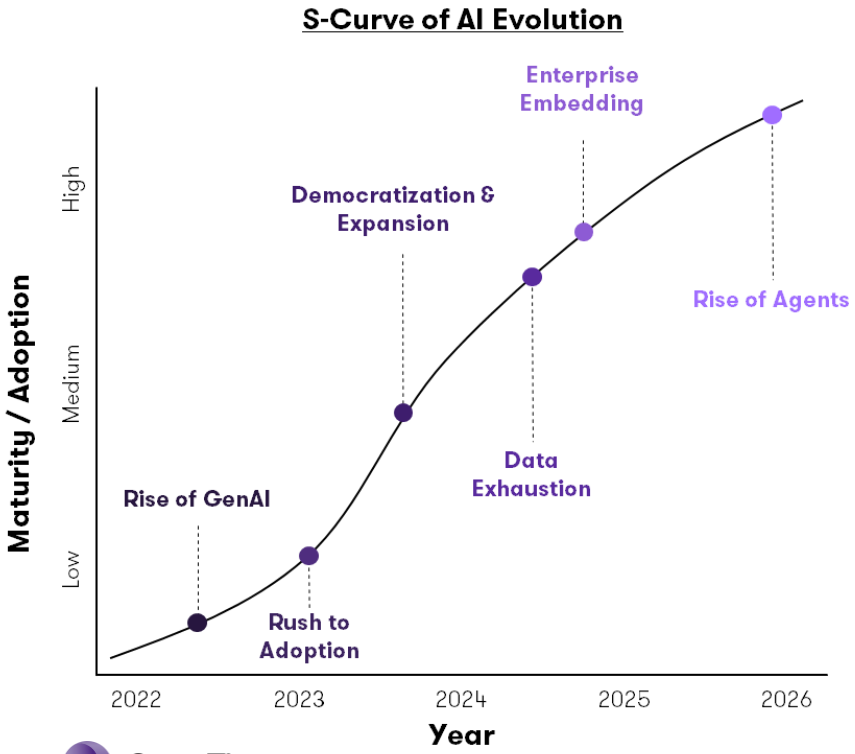

- **Rise of GenAI:** Breakthroughs in large language models (LLMs) such as ChatGPT triggered the generative AI boom, enabling scalable content creation, translation, summarization, and multimodal applications.
- **Rush to Adoption:** Companies raced to deploy AI for cost savings and new experiences-often at the expense of foundational stability, leading to mounting tech debt.
- **Democratization & Expansion:** APIs, open-source models, and low/no code tools unlocked rapid adoption, empowering teams to build and scale AI with fewer barriers.
- **Data Exhaustion:** With most quality public data already consumed, firms are now racing to secure proprietary, domain-specific data to fuel next wave.
- **Embedding AI into the Enterprise:** Companies are moving from isolated pilots to end-to-end AI integration-transforming workflows while grappling with unclear ROI.
- **Rise of Agents and SLMs:** Focus is shifting to smaller, domain-specific models and autonomous agents-bringing AI closer to decisions with precision and judgment.

Grant Thornton

Why now?

The case for action is not hype; it is a control design problem created by autonomy, tool use, and multi-step behavior.

82%

of executives plan to adopt AI agents within the next one to three years

64%

of audit teams are exploring or considering AI agents within 12 months

70%

of audit and assurance professionals expect to need materially more AI skills within a year

Traditional AI risk	GenAI risk	Agentic AI risk
(output quality)	(prompt + content)	behavior + action + coordination

Implication for Internal Audit

- Traditional audit asks whether the output is correct.
- Agentic audit must also ask whether the system stayed inside approved permissions, guardrails, and escalation paths.
- That changes advisory strategy, assurance scope, and required evidence.

What makes agentic AI different

GenAI answers. Agentic AI acts.

Perceive

Reads context, data, and goals

→

Plan

Breaks a goal into steps

→

Act

Calls tools and executes — with limited human input

Co-pilot

You prompt — it responds. A tool you invoke.

Agent

It initiates and completes tasks. An actor on the team.

*The shift from **tool you invoke** to **actor that initiates** is exactly what creates new governance needs.*

Grant Thornton

Identify where control exists in AI-enabled systems, particularly within the execution layer

(permissions, approvals, and logging).

What makes up Agentic AI

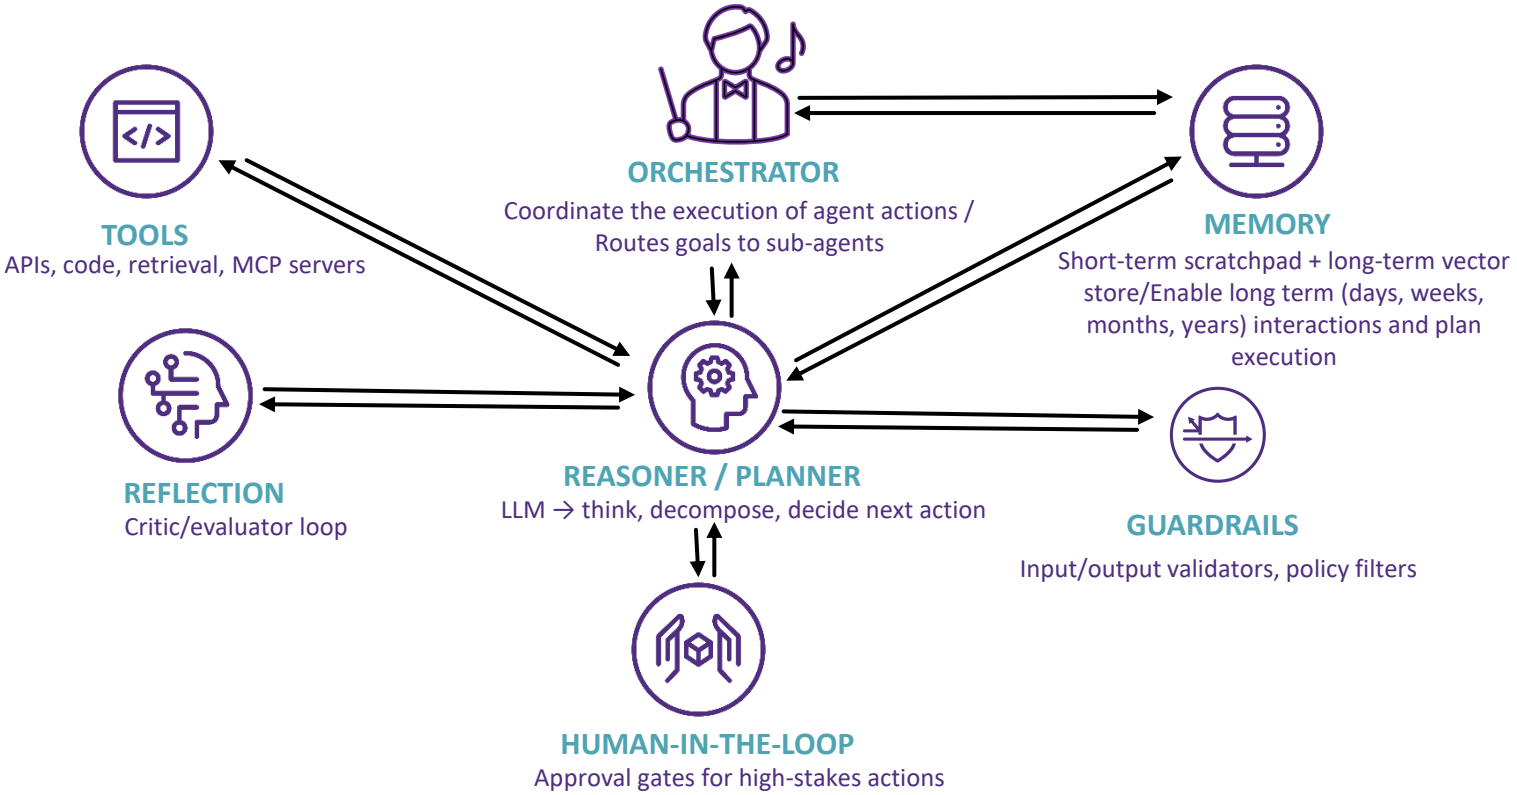

The harness is where controls live — it turns model intelligence into auditable, bounded enterprise behavior

If the model is the “brain,” the harness is the control layer that determines what the agent can see, do, approve, log, and recover from.

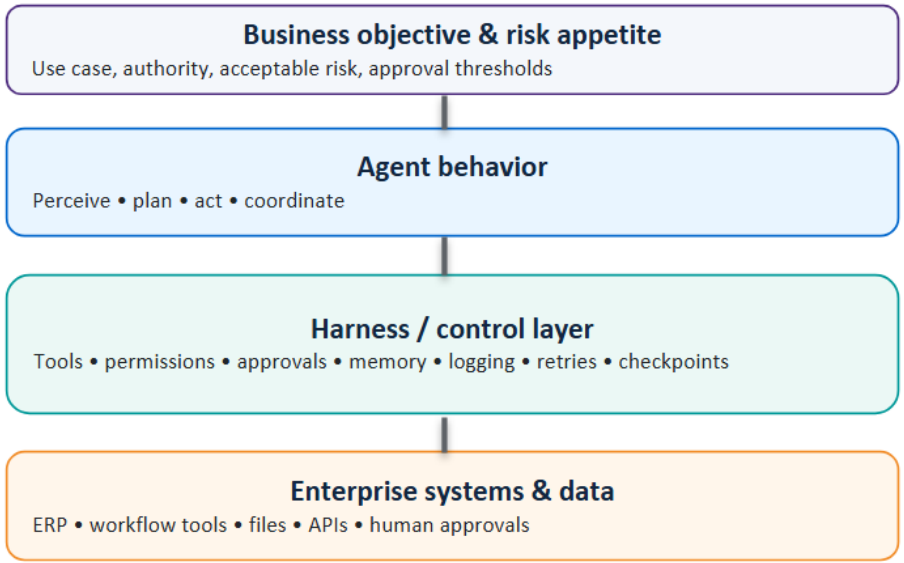

Control Environment

Ownership, accountability, escalation

Risk Assessment

Autonomy, blast radius, use-case triage

Control Activities

Permissions, limits, approvals, HITL

Information & Communication

Action logging, explainability, traceability

Monitoring

Audit the harness if you want assurance over agent behavior.

Grant Thornton

Recognize how audit approaches must adapt based on levels of AI autonomy and the associated business impact (risk) of AI-enabled activities.

AI system autonomy and business impact drives the audit approach — the wrong approach applied to the wrong level will either miss risk or over-control innovation

Co-pilot / assistant	Semi-autonomous agent	Autonomous agent	Swarm / multi-agent system
Waits for prompts; no independent action	Recommends or executes with human approval	Acts within pre-set permissions and limits	Multiple agents coordinate and hand off work
Primary audit challenge Primarily output quality and over-reliance	**Primary audit challenge** Human override quality and explainability	**Primary audit challenge** Control architecture, bounded autonomy, failure handling	**Primary audit challenge** Emergent behavior, information flow, and authority conflicts

Management needs a common language for autonomy; Internal Audit needs a tiered response model.

Audit approach progression

• Co-pilot: validate outputs. • Semi-autonomous: validate outputs + human decisions. • Autonomous: audit permissions, thresholds, escalation logic, and logs. • Swarm: audit system behavior, coordination protocols, and end-to-end traceability.

Agentic adoption tiers from a deployment lens

	Tier	Core Aspects	Examples & Use Case
Tier 0	Vendor embedded assistant	Fully vendor-controlled. Organizations consume it, not build it.	Examples: M365 Copilot, Salesforce Einstein Use case: productivity drafts, internal summarization with no sensitive data
Tier 1	Citizen developer agent	Low-code/no-code platform. User configures flows and prompts, not code. Actions on real org data.	Examples: Power Automate + AI Builder, Copilot Studio, Zapier Central AI, Google AppSheet AI Use case: internal productivity; no sensitive data; no autonomous actions; approved tools/patterns
Tier 2	Custom in-house agent	You built it. You control identity, tools, and boundaries.	Examples: Amazon Bedrock Agents, In-house (retrieval augmented generation(RAG) pipelines, Custom agents Use case: agents autonomously handle routine chores like monitoring specialized databases, custom integrations
Tier 3	Externally extended agent	Connects to external tools/services across trust boundaries.	Examples: Agents with Model Context Protocol servers, agents calling third-party APIs Use case: initiate updates to customer records (e.g. refunds), material financial/operational impact, sensitive data, or agentic actions
Tier 4	Multi-agent orchestration/swarm of agents	Multiple agents coordinate within your organization.	Examples: Crew AI workflows, AutoGen multi agent teams,LangGraph multi-agent graphs Use case: form filling & data extraction pipelines, autonomous,multi-agent system and high-blast-radius actions

Grant Thornton

Redefine Internal Audit's role in AI-driven environments by translating the concept of a "trusted advisor" into earlier engagement, behavior-level assurance, and delivering actionable insight where the business lacks visibility.

Internal Audit delivers value by assisting organizations in transforming ideas into results and offering an independent perspective that reveals potential risks within opportunities.

Rapid innovation

Efficiency is no longer optional

71% **of companies already use GenAI in at least one function**
McKinsey, The state of AI

Rising costs, regulatory complexity, and margin compression are forcing leaders to do more with less – and AI is the only lever that scales profitability

Technology maturity

The ecosystem is ready; velocity is the gap

$644B **global AI spend expected by 2028, up from $307B is 2025**
IDC, Unlock the Future of AI

Foundational models, secure cloud, and enterprise-grade copilots have de-risked adoption – what's missing now is disciplined acceleration

Audit Focus on IT & Cyber

Work is being redesigned around new technologies and specialist audits.

80% **Cited focusing on audit towards IT of which 66% is focused on Cybersecurity**
The IIA.org, March 2024

AI agents are moving from pilot projects to production, augmenting human capability and re-wiring workflows for speed and accuracy

Audit Plans for FY26

Audit plans are continuing to focus on technology with

20% **of audit plans as compared to operational auditing (17%), compliance and financial reporting both at (14%)**
Grant Thornton, Digital Sense Survey 2025

As more investments are made into AI and technology initiatives, there is limited clarity on how it connects to outcomes, ROI, Safety and Security – the whitespace GT helps close through its advisory services.

Grant Thornton

Common failure points – governance lens

AI adoption is outpacing governance practices, especially as agents and citizen development decentralize AI creation.

Review speed bottleneck

Central review queues cannot keep up with business-led pilots, copilots, and agent development

Lack of evidence validation

Questionnaires and attestations often replace evidence-based validation and operating telemetry

Citizen development sprawl

Non-technical builders create AI workflows without consistent design, testing, logging, or release discipline

Third-party AI opacity

Vendor AI is embedded in SaaS, black-box, updated frequently, and difficult to test independently

Model and data drift

Behavior changes as data, prompts, retrieval sources, tools, and model versions change

Autonomy and permissions

Agents can chain tasks, use tools, access data, and act as non-human identities with real business impact

Regulatory uncertainty

AI, privacy, cyber, employment, consumer, and industry rules are overlapping and evolving

Incident readiness gap

Organizations lack monitoring, escalation thresholds, and tested kill-switch procedures for AI harms

Implication: governance must be embedded, risk-informed, and continuously monitored; not managed only through point-in-time reviews.

Internal Audit as a trusted advisor

The value of advisory is not generic governance commentary; it is forcing clarity on ownership, autonomy, and failure modes early.

What management is trying to do	How Internal Audit adds value	What good looks like
Select a use case	Challenge whether the use case starts with high-volume, low-judgment work or jumps too quickly to high-stakes autonomy	Pilot starts in a bounded workflow with observable evidence and a defined owner
Set the degree of autonomy	Push for a clear definition of what the agent may decide, what requires human approval, and what is prohibited	Autonomy thresholds, approval triggers, and escalation paths are documented before build
Design the operating model	Translate business intent into permissions, logging, and exception handling requirements	Harness controls are specified up front: least privilege, logging, checkpoints, and override paths
Prepare for failure	Force the team to define “wrong” before launch — data errors, tool failures, prompt injection, or runaway loops	Management has clear kill-switch, incident, and rollback decisions tied to the use case

Trusted advisory = translating strategy into control architecture before the agent touches production systems.

Assurance must progress from output validation to architecture testing to system-behavior analysis as autonomy increases

The same audit approach cannot be applied consistently across semi-autonomous agents, autonomous agents, and swarms.

Autonomy	Primary assurance question	What Internal Audit tests	Evidence expected
Semi-autonomous	Did the human stay meaningfully in the loop?	Recommendation quality, override behavior, rationale visibility, evidence of review	Output logs, approval records, user guidance, exception handling
Autonomous	Did the system operate inside approved permissions and thresholds?	Permissions, tool calls, escalation logic, incident handling, retry and stop conditions	Step-level logs, execution traces, controls over write-back and exceptions
Swarm / multi-agent	Did the system coordinate predictably and protect information across agents?	Task delegation, authority conflicts, information flow, trace join-up across agents	Cross-agent traces, communication logs, scenario test results, segregation of duties evidence

For autonomous systems, “correct output” is not enough; the path taken to get there becomes auditable.

Correct output is not enough

Scenario

- A reconciliation agent correctly identifies unmatched transactions and prepares the right journal-entry exception list.
- But to get there, it accesses a restricted dataset, retries against a failing API until thresholds are exceeded, and triggers an automated workflow without the intended approval step.
- The business user sees the right exception list. Internal Audit should still conclude behavior fell outside approved boundaries.

Outcome test
"Was the answer correct?"

Path step 1: Accessed restricted data

Path step 2: Exceeded retry threshold

Path step 3: Bypassed intended approval

Audit conclusion: validate the outcome, but assess the trajectory.

Grant Thornton

Grant Thornton

Define the future role of Internal Audit in AI-driven environments by distinguishing between “audit with AI” and “audit of AI,” and explaining the shift from output validation to control architecture and system-behavior assurance.

IA now operates in two modes simultaneously

Internal audit now occupies a dual role — deploying AI to expand coverage while simultaneously providing assurance over the organization's AI systems — and these responsibilities reinforce each other.

Audit with AI

Efficiency Lever

- Broader coverage and speed — agents test full populations where staff once sampled
- Continuous, real-time assurance replaces periodic, point-in-time reviews
- More auditor time freed for judgment, risk dialogue, and strategic counsel

AND

Audit the AI

Assurance Obligation

- Governance and accountability of AI systems — who owns each agent and its outputs
- Model behavior and data lineage — tracing how conclusions were reached
- Controls over autonomous action — ensuring agents operate within sanctioned boundaries

Deloitte's 2026 Internal Audit Hot Topics identifies governing agentic AI systems as the single mega hot topic for the profession — the defining challenge every audit plan must address. [6]

Audit with GenAI/AI/Agentic AI

	Planning & Risk Assessment	Fieldwork	Reporting Outputs	Insight & Oversight Activities
Execution Efficiency	Synthesize risks Draft audit plan Draft scoping memo Risk brainstorming Policy/SOP summarization & risk extraction	Draft narratives Draft risk and controls Document request list creation	Draft findings memo Draft reports Draft audit summary of results (stakeholder communications)	Controls rationalization
Audit Quality & Consistency	RACM creation and analysis Draft plan RACM analysis against relevant framework Analysis of narratives and flowcharts for control gaps	Control attribute definition Assist with sampling strategy Analyze narratives amd flowcharts for control gaps Assist with root cause analysis Workpaper creation Unstructured data testing	Standardization checks for 5C structure and conclusions Analysis of reviewer comments and rework patterns Traceability mapping finding → control → evidence → conclusion Automated assessment of:Missing documentation Closing meeting & report drafting	Issue tracker summaries Post-mortem/Process improvements synthesis
Continuous Audit Intelligence & Risk Monitoring	Identification of leading indicators of control failure Heatmaps identifying high-risk processes, entities, or geographies Prioritization based on risk trajectory not historical results GenAI-assisted development of targeted audit scopes	Automatically generate process flowcharts from walkthrough meeting transcripts Control testing with agentic AI Automated testing of 100% of transaction volume Full autonomous auditing model	KRI / KCI trend dashboards with narrative insights Repeat findings and systemic control failure identification Dynamic audit committee briefs Enterprise-wide control weakness patterns Predictive indicators based on historical issue patterns	Control and risk signal interpretation (KRI/KCI) Data-driven recommendations for automation Issue tracker insights and aggregation SOX compliance agent as part of verification swarm of digital specialists

*Agentic AI

Grant Thornton

From output validation to architecture testing and system-behavior analysis

Agent System
(Co-pilot → Agents → Multi-agent Swarms)

Execution Path
Perceive → Plan → Act → Interact

Control Layer (Harness)
Permissions | Approvals | Logging | Limits

Using AI expands assurance delivery, auditing AI expands assurance scope. Both require new audit capability.

Capacity Implication

- Build capability to analyze **execution logs** and reconstruct decision-making.
- Build capability to evaluate **control architecture** and whether agents operate within predefined constraints.
- Build capability to review **harness / control layer** design, including permissions, approvals, and logging.
- Build capability to assess **end-to-end traceability** and behavior across autonomous and multi-agent systems.

Agentic systems break outcome-based assurance. Internal Audit's role as a trusted advisor shifts upstream into system design, while assurance evolves from output validation to control architecture and behavior analysis—or remains a reactive validator of outcomes.

Open for Q & A

Contact us

Matthew Cassidy

Partner,
Internal Audit & SOX
Grant Thornton Advisors LLC
matt.cassidy@us.gt.com

Daryl Ray

Exp Manager,
Internal Audit & SOX
Grant Thornton Advisors LLC
daryl.ray@us.gt.com

Visit us online

www.gt.com/ai

linkedin.com/company/grant-thornton-us/

This Grant Thornton Advisors LLC presentation is not a comprehensive analysis of the subject matters covered and may include proposed guidance that is subject to change before it is issued in final form. All relevant facts and circumstances, including the pertinent authoritative literature, need to be considered to arrive at conclusions that comply with matters addressed in this presentation. The views and interpretations expressed in the presentation are those of the presenters, and the presentation is not intended to provide advice or guidance with respect to the matters covered. It is not, and should not be construed as, accounting, legal, tax, or professional advice provided by Grant Thornton Advisors LLC.

Grant Thornton Advisors LLC and its subsidiary entities are not licensed CPA firms. For additional information on matters covered in this presentation, contact a Grant Thornton Advisors LLC professional.

This Grant Thornton Advisors LLC content provides information and comments on current issues and developments. It is not a comprehensive analysis of the subject matter covered. It is not, and should not be construed as, accounting, legal, tax, or professional advice provided by Grant Thornton Advisors LLC. All relevant facts and circumstances, including the pertinent authoritative literature, need to be considered to arrive at conclusions that comply with matters addressed in this content.

Grant Thornton Advisors LLC and its subsidiary entities are not licensed CPA firms.

For additional information on topics covered in this content, contact a Grant Thornton Advisors LLC professional.

“Grant Thornton” refers to the brand name under which the Grant Thornton member firms provide services to their clients and/or refers to one or more member firms, as the context requires.

Grant Thornton LLP and Grant Thornton Advisors LLC (and their respective subsidiary entities) practice as an alternative practice structure in accordance with the AICPA Code of Professional Conduct and applicable law, regulations and professional standards. Grant Thornton LLP is a licensed independent CPA firm that provides attest services to its clients, and Grant Thornton Advisors LLC and its subsidiary entities provide tax and business consulting services to their clients. Grant Thornton Advisors LLC and its subsidiary entities are not licensed CPA firms.

Grant Thornton International Limited (GTIL) and the member firms, including Grant Thornton LLP and Grant Thornton Advisors LLC, are not a worldwide partnership. GTIL and each member firm are separate legal entities. Services are delivered by the member firms, GTIL does not provide services to clients. GTIL and its member firms are not agents of, and do not obligate, one another and are not liable for one another’s acts or omissions.