

2026 REPORT

AI Impact Survey

Understanding the AI proof gap and why AI isn't delivering the performance leaders expected.



Most organizations are scaling AI they cannot explain, measure or defend.

This is the AI proof gap.

78%

of organizations lack strong confidence that they could pass an independent AI governance audit. This report explains why — and what the organizations pulling ahead are doing differently.

The bottom line for leaders

Governance is about performance, not just compliance.

Organizations that claim fully integrated AI are almost four times more likely to report revenue growth than those still piloting. The difference proves to be governance, not technology.

Strategy separates the leaders. And most organizations still don't have one.

51% of all organizations name strategy as the biggest driver of AI ROI — above talent, tech stack and governance combined. Yet only 22% have a fully developed and implemented enterprise AI strategy in their operations function. The organizations pulling ahead turn their vision into a fully designed plan for applying AI to specific processes and workflows.

Your C-suite views reality differently.

CIOs and CTOs are five times more likely than COOs to say the workforce is ready for AI. That misalignment is undermining the outcomes their organizations are trying to achieve. Until leadership aligns, the gap will widen.

INTRODUCTION

Executives are scaling AI. They are not governing it.

The AI proof gap is showing up in revenue, innovation velocity and the measurable distance between organizations that are scaling AI with discipline and those that are not.

 **78% of business executives in Grant Thornton's 2026 AI Impact Survey lack strong confidence that they could pass an independent AI governance audit within 90 days.**

Most have approved investments and launched initiatives, ranging from piloting to scaling to full AI integration. But they can't explain, measure or defend their AI, and they're not prepared to handle the fallout if it fails.

Grant Thornton's 2026 AI Impact Survey, conducted through live interviews with 950 business leaders across 10 industries from Feb. 13 to March 18, 2026, reveals a widening gap between organizations that can prove their AI is generating measurable results and those that cannot.

The data surfaces a striking contradiction. Governance and compliance failures are the leading cause of AI underperformance, cited by 46% of business executives. Yet only 11% identify risk and compliance as the function that needs the most focus to meet their AI ambitions. This gap between knowing what is failing and not prioritizing the solution is itself a dimension of the proof gap.

For most, a lack of C-suite alignment is slowing progress and escalating risks. Inside many organizations, COOs overseeing AI-affected operations are discovering governance gaps that CFOs are not funding and that CIOs and CTOs are not surfacing. The result: AI scaling without anyone accountable for what it produces.

This report explores four critical dimensions of the AI proof gap — governance, strategic focus, workforce readiness and the rising risks of agentic AI — before revealing the measurable performance differences between organizations that are closing the gap and those falling behind. Together, these insights show where real value is being created and how leaders can capture it.

The organizations pulling ahead have built governance that gives their leaders the confidence to scale AI decisively. The rest are inheriting risks they cannot see and outcomes they cannot prove. Where the gap is widening, how it is affecting performance and what it takes to close the performance gap are explored next.

"Across the organizations we work with, what we consistently see is that AI deployment has outpaced the infrastructure to defend it. Leaders who have invested in governance aren't moving slower — they're moving faster because they have the confidence to scale. The ones who haven't built it yet are one incident away from a much harder conversation."

Tom Puthiyamadam,
Managing Partner, Advisory Services,
Grant Thornton Advisors LLC

Your AI investment was approved. No one owns the outcome.

75%

of boards have approved
major AI investments

— HOWEVER —

48%

of boards have
not set governance
expectations for AI

 **Boards are signing off on AI, but many are not asking what happens if something goes wrong.**

Three in four boards have approved major AI investments, but 48% have not set AI governance expectations and 46% have not integrated AI risk and opportunity into ongoing board or committee oversight. This shows that approvals are being made without the mechanisms to know whether what they approved is working or failing.

That gap runs deeper than board oversight. Most governance models were not built for the volume of AI use cases organizations are now deploying. Centralized review bodies get overwhelmed as use cases multiply, creating bottlenecks that slow the business without actually reducing risk.


AI governance is the enterprise system that determines who owns AI decisions, who can demonstrate they are working and who is accountable when they are not. It is the operational infrastructure that determines whether AI can scale without becoming a liability.

🌀 The proof gap is real and it is measurable. The question is: what separates the organizations that can prove their AI works from those that cannot. The answer, revealed consistently across the survey data, is governance. Not governance as most organizations practice it, but governance built as a performance system.

Governance gaps lead to AI failure

Governance or compliance barriers are the leading cause of AI underperformance.

Q: Choose up to three factors that have contributed to AI project underperformance or failure in your organization.



Source: Grant Thornton's 2026 AI Impact Survey. Base of 950.

Governance or compliance barriers dominate at 46%, the single largest factor by a wide margin. But the finding at the bottom of the chart may be the most telling: Only 8% of respondents selected "AI is not underperforming." That means the majority of business leaders surveyed implicitly acknowledge that AI is underperforming in their organization.

Most organizations are running AI they cannot fully account for.

The gap between what organizations are deploying and what they can prove works is growing with every initiative that scales without a governance structure. In return, AI initiatives are stalling across organizations that have the tools, the budget and the mandate to scale. Without strong

governance, piloting and scaling produce activity, not outcomes. Organizations are moving through discovery and deployment without building the ability to show that AI is working safely, defensibly and at the scale the business requires. Each ungoverned initiative does not just create one gap. It creates a gap that makes the next initiative harder to govern, harder to measure and harder to defend. The proof gap does not grow linearly. It compounds.

Governance aligns with AI maturity

Q: How confident are you that you could pass an independent audit of AI governance and controls (with documented evidence) within 90 days?

% of respondents answering: “Very confident” (Controls are defined and operating and evidence is centralized)



Source: Grant Thornton's 2026 AI Impact Survey. Base of 950.

None of the 28 “early AI exploration” stage respondents were “very confident” they could pass an independent AI governance audit. Proof at the earliest stages is not low. It is nonexistent. Organizations do not drift into governance confidence. They build it deliberately. The gap between piloting and fully integrated is tenfold.

The organizations managing this well have stopped trying to govern everything from the center. They set policy and risk criteria centrally and delegate assessments to trained reviewers at divisional or regional levels. Not every use case carries the same risk. Directing scrutiny where it matters and letting lower-risk decisions move quickly allows organizations to scale AI without increasing their exposure.

Most have not arrived there yet. Only 40% are well-prepared to handle the privacy and security challenges that AI creates as multinational privacy risks multiply and increase complexity. The regulatory dimension compounds the governance gap, and the specifics vary by industry.


Among organizations still piloting AI, just 7% are very confident their organization could pass an independent AI governance audit today. Among those claiming fully integrated AI, 74% are. The organizations at the top built governance as the foundation for scale. The ones still piloting are not behind on technology. They are behind on the infrastructure to prove their technology works. That is the proof gap.

Industry lens: Controls confidence lacking

PE, manufacturing and construction/real estate struggle with AI governance.

Q: How confident are you that your organization could pass an independent audit of AI governance and controls (with documented evidence) in the next 90 days?

Chart shows percentage of respondents who answered "Very confident — Controls are defined and operating, and evidence is centralized."



Source: Grant Thornton's 2026 AI Impact Survey. Base of 950.

The proof gap is not uniform across industries. Technology and telecommunications lead at 42% confidence, while private equity sits at 9%. That spread matters because it shows that industry context shapes governance readiness.

Strategy drives ROI. Three in four haven't built one.

51%

say strategy is the
biggest driver of AI ROI

HOWEVER

79%

of operations
respondents do not have
a fully developed and
implemented AI strategy

More than half of executives say strategy is the biggest driver of AI ROI. Fewer than one in four have actually built one. Strategically, organizations are succeeding on breadth — more pilots, use cases and functions touched by AI. They are failing on depth.

That's because many are motivated by the fear of falling behind instead of a clear view of where AI creates value. Leaders identified competitor moves as the biggest external pressure driving adoption, and that produces breadth without depth while bypassing the feedback loops needed to know whether any of it is working.



**Depth turns a roadmap
into outcomes.**

Building measurement targets and governance infrastructure enables teams to move faster with confidence. That means consistent ROI measurement across initiatives, feedback loops that inform where the next investment should go and the courage to exit experiments that are not delivering. It also means starting where the evidence is easiest to build.

Functions with lower data sensitivity such as marketing, procurement and internal operations often offer the fastest path to demonstrable returns. Those early results justify deeper investment. Depth pays dividends that are portable to other aspects of the organization.


Leaders also need a clear view of their industry's specific sources of competitive differentiation to define where AI should create value. Without that clarity, organizations default to generic AI enablement — chasing tools and benchmarks that apply to broad swaths of employees or generic functions instead of defining discrete problems that AI should solve in a specific context to create strategic advantage.

The organizations that built a strategy can measure what is working and prove what is not. The ones still operating on breadth are spending more, launching more and able to defend less. Strategy without measurement is not strategy. It is activity with a roadmap.

In search of a strategy

Most leaders haven't built or implemented a comprehensive AI strategy.

Q: Does your organization have a formal enterprise AI strategy?



Note: Figures do not add to 100% due to rounding

Source: Grant Thornton's 2026 AI Impact Survey.
Operations business leader respondents. N=390

Only 22% of organizations have a fully developed and implemented AI strategy while 53% are currently in development — building the plane while flying it.

AI is speeding ahead. The workforce isn't ready.

AI performance isn't matching aspirations because workforce readiness is lagging. Training is disconnected from workflows and tasks, resulting in confusion for employees. Leadership misalignment is to blame.

Only 12% of leaders say their workforce is fully ready to adopt AI. Most (81%) say their workforce is "fairly" or "mostly" ready. That distance between 12% and 81% is the workforce's defining tension. "Mostly ready" is not close enough. It means organizations cannot demonstrate that their workforce can execute AI-enabled work to a measurable standard. It is the comfort of proximity without the substance of arrival.

The problem is misalignment.

Within the C-suite, CIOs and CTOs (at 39%) are five times more likely than COOs (at 7%) to say their workforce is fully ready to adopt AI. That gap reflects two different views, each with substantial organizational pull, moving in different directions within the same organization. CIOs and CTOs are likely to measure adoption at the system level: usage statistics, tool access, platform rollout. COOs are likely to measure at the operational level: how people implement the technology to advance business goals. The

34%

cite training as the most underfunded AI investment area, ranking it first among all categories.

distance between those two perspectives leads to a disconnect on moving the workforce forward toward organizational goals and optimizing AI performance.

The most common response to the workforce gap is more training. But awareness training merely moves people toward the same starting line. It creates a baseline of familiarity. It does not create capability. Too many organizations stop there, treating AI implementation like any other IT project: provide access to tools, run enablement sessions, check the box. But AI adoption does

not follow the pattern of prior technology rollouts. It requires changing habits, decision-making and workflows at the role level so AI becomes part of how work gets done, not just another tool employees have access to.

The people closest to AI in daily operations need the most support to implement it, yet they are getting the least. Middle managers are diminishing in numbers while the workload for those remaining has accelerated rapidly. The layer of the organization most responsible for making AI work in practice is simultaneously being thinned and overloaded.

Combined, frontline employees (37%) and middle managers (30%) represent roughly two-thirds of where organizations say support is most needed. The people closest to AI in daily operations, translating leadership's AI ambitions into actual workflows, are the ones least equipped to do it. Meanwhile, senior leadership and functional leadership, combined, sit in single digits. While the people making AI decisions feel adequately supported, the people living with those decisions do not. Only 3% say they have not seen significant adoption challenges, meaning 97% acknowledge the problem exists.

The organizations with the strongest workforce readiness didn't get here through more training, but through leadership that redefined how work gets done. Their leaders defined where AI should be applied, what decisions it should influence and what outcomes employees are accountable for. They are equipping their workforce to use AI tools within specific, defined workflows with clear rules about usage and accountability. The difference is not access to better tools or more training budget. It is leaders who redesign work rather than merely providing employees with access to technology.

The workforce gap will not close with more training dollars or broader tool access. It will close when leadership aligns on what readiness means, defines it at the role level and redesigns work to meet it. Until then, organizations will keep training their workforce for a race without a finish line.

Support needed on the front line

Employees charged with execution need help implementing AI.

Q: Which part of the organization currently requires the most support to effectively adopt AI-enabled ways of working?



Source: Grant Thornton's 2026 AI Impact Survey. Base of 950.

Agentic AI is accelerating. Most aren't prepared for its failure.

ONLY **1 in 5**

organizations have a tested
response path for AI failure

— YET —

74%

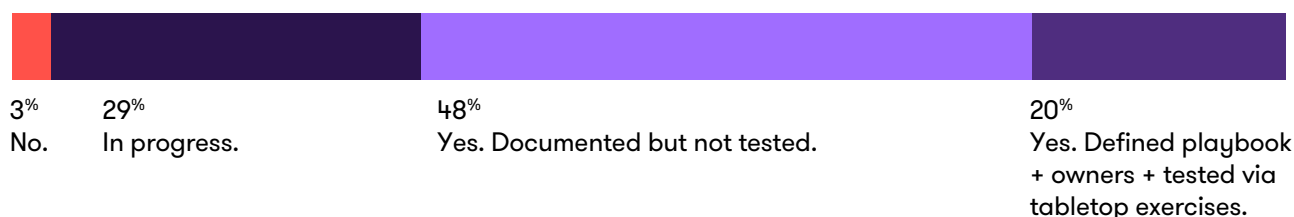
of business leaders are
giving agentic AI access to
their data and processes

Nearly 3 in 4 organizations are piloting, scaling or running autonomous AI. Despite this, just 1 in 5 has tested a response plan for when something goes wrong. Most organizations are proceeding cautiously with how much autonomy they grant: Only 5% permit agents to execute high-stakes decisions without human review and 60% limit agents to moderate-risk task automation. That caution is rational. What is not rational is that even at those limited levels of autonomy, the governance infrastructure has not kept pace.

AI incident playbooks mostly untrusted

Most haven't practiced an AI-specific response plan. Some don't have a plan at all.

Q: Do you have an AI-specific incident response playbook (bias, hallucinations, harmful outputs, model drift, data leakage, misuse)?



Source: Grant Thornton's 2026 AI Impact Survey. Base of 950.

The data on preparedness is unambiguous. Only 20% of organizations have a fully defined AI incident response playbook with assigned owners that has been tested through tabletop exercises. The rest fall across a spectrum of unpreparedness: 48% have a documented playbook they have never tested, 29% have a playbook in progress and 3% have no playbook at all.

The dominant slice is not “no.” It is “yes, documented but not tested.” Nearly half of organizations have a playbook sitting in a drawer that has never been pressure-tested against a real scenario. A plan you have not tested is a hypothesis, not a plan. It gives leadership the comfort of feeling prepared without the reality of being prepared. That false confidence may be more dangerous than having no plan at all, because it removes the urgency to act.

 **48% of organizations have a documented AI incident playbook that they have never tested. A plan you have not tested is a hypothesis, not a plan.**

The absence of tested playbooks is the visible symptom. The underlying condition is that most organizations have no governance structure designed for autonomous systems. Building a playbook today does not solve the problem if there is no mechanism to update that playbook as agents evolve, as use cases expand and as the risk landscape shifts.

The governance most organizations have built was designed for a world where humans make decisions with AI assistance. In that world, static policies and periodic reviews are sufficient because there is a human checkpoint at every decision point. Agentic AI removes that checkpoint. When agents operate autonomously, making decisions and taking actions without waiting for human review, governance must operate at the same speed and with the same continuity as the agents themselves.

Governance for agentic AI must pivot from static documentation to continuous oversight. That means monitoring agent behavior in real time, detecting deviations from expected parameters and adjusting controls as systems evolve. It means governance that operates as an ongoing function, not a one-time exercise. Organizations that build this capability now will scale agentic AI without increasing their exposure. They will be able to contain incidents quickly and reconstruct how an outcome was produced and where accountability sits.

The shift to continuous oversight does not require building from scratch. Most organizations already have business continuity planning and incident response protocols. Employees know where to report, who to call and how to escalate when triggering events occur. That existing infrastructure needs to be extended to contemplate AI-specific failure modes: model drift, hallucination, unauthorized data access, accidental data leakage and cascading workflow errors. Rapid containment, measured risk management and clear accountability for AI do not require entirely new systems. They require adapting what already exists for a new category of risk.

Policies written before agents were designed and deployed will not cover what happens when they run at scale.

Within the C-suite, the misalignment this report has tracked through workforce readiness reaches its most consequential form around agentic AI. COOs (54%) are almost three times more concerned than CIOs and CTOs (20%) about regulatory and compliance uncertainty related to agentic AI. This is no longer a disagreement about whether the workforce is ready. It is a disagreement about how exposed the organization is to regulatory risk from autonomous systems.

When two leaders who share accountability for AI deployment disagree by this margin on regulatory exposure, the organization cannot produce a coherent account of its risk posture. As agentic AI deployment scales, that gap in alignment around accountability and risk tolerance makes it harder to maintain control, harder to respond to regulatory scrutiny and harder to prove the organization has its autonomous systems governed.

The gaps in governance, strategy and workforce readiness that this report has documented are consequential when humans are making AI-assisted decisions. They become dangerous when AI systems begin operating on their own. Each new agent that accesses organizational data and processes without continuous oversight adds to the organization's unmonitored exposure. Each day that governance operates on a static model while agents operate in real time widens the gap between what the organization is running and what it can account for. The agentic proof gap does not hold steady while organizations deliberate. It accelerates.


The question is no longer whether your organization will experience an agentic AI failure. It is whether you have built the infrastructure to identify it, contain it, remediate it and explain it. Most have not.

Industry lens: AI incident response

Manufacturing and PE are lagging.

Q: Do you have an AI-specific incident response playbook (bias, hallucinations, harmful outputs, model drift, data leakage, misuse)?

Chart shows percentage who answered "Yes — defined playbook + owners + tested via tabletop exercises."



Source: Grant Thornton's 2026 AI Impact Survey. Base of 950.

The spread between technology and telecommunications (39%) and manufacturing (7%) reveals that preparedness is wildly uneven across entire industries. Manufacturing sits at the bottom despite facing some of the most tangible consequences of AI failure: operational downtime, safety incidents and physical process disruption. In an industry where the cost of an uncontained AI failure can be measured in halted production lines and worker safety, 93% have not pressure-tested their response.

The gap is measurable. Governance is what separates the leaders.



59%

of organizations with
fully integrated AI report
accelerated innovation

HOWEVER


24%

of organizations piloting
AI report accelerated
innovation

The four dimensions of the proof gap — governance infrastructure, strategic discipline, workforce readiness and agentic AI preparedness — are not separate problems with separate consequences. They are interconnected gaps that produce a single, measurable outcome; organizations that have closed them are outperforming those that have not across every metric the survey measured.

The beginning of this report showed that organizations with fully integrated AI are four times more likely to report revenue growth than those still piloting. What follows is the full picture. The data is clear. The almost four times revenue gap is not an isolated advantage. It is one dimension of a structural performance separation that is consistent across revenue, efficiency, innovation and output quality.

AI integration delivers outsized benefits



Organizations with fully integrated AI are reporting gains across every measured outcome. Efficiency leads at 81%, meaning more than four in five fully integrated organizations report measurable efficiency improvements from AI. Revenue growth reaches 58%, output quality 64% and innovation 59%. The pattern is what matters: fully integrated organizations are not excelling in one area and lagging in others. They are outperforming across the board. That consistency is the signature of governance as a performance system. These organizations built the infrastructure to make AI provable and defensible across everything it touches, and the results followed.

At the other end of the table, the story is equally clear. Organizations still piloting AI report revenue growth at 15%, innovation at 24% and output quality at 35%. For revenue, that means 85% of piloting-stage organizations are seeing no measurable revenue impact from AI at all. The only metric where piloting organizations show a majority benefit is efficiency at 58%, and even there the gap to fully integrated is 23 points. Piloting-stage organizations are not simply behind their competitors. For most, the AI investment is producing minimal measurable returns. The proof gap at the piloting stage is about organizations continuing to make large investments and being unable to demonstrate real outcomes.

The survey reveals something the top-line comparisons obscure: the leap from scaling to fully integrated is where the most dramatic gains appear for most of the benefits. Innovation nearly doubles from 29% to 59%. Efficiency moves from 63% to 81%. The acceleration at the fully integrated stage suggests that the final step — from scaling AI with some governance to operating AI with full governance integration — is where the compounding advantage becomes most visible. Organizations do not drift into fully integrated status. They build the governance infrastructure that makes the leap possible. That infrastructure is what produces the tenfold gap in audit confidence.

These numbers reflect organizations at different points on the same journey. The ones at the top did not get there by accident, and the ones at the bottom are not there because of technology limitations. The gap between the leaders and the rest is the cost of the proof gap, measured in outcomes. It is driven by governance. It compounds through strategy, workforce and agentic readiness. And it is visible at every stage of maturity, from the near-zero returns at piloting to the structural advantage at full integration. The distance is already measurable. It is already widening. And it will not close without deliberate accountability.

Defensible AI delivers results — but it takes discipline to build

The proof gap is an accountability problem. Boards approved investments without setting governance expectations. Leadership deployed AI without defining who owns the outcomes. Organizations scaled without building the infrastructure to prove any of it works.

Focusing on these priorities can help leaders close the AI proof gap:

1 **Build governance like a performance system, not a policy document**

Organizations with fully integrated AI are 10 times more likely to pass an independent governance audit than those still piloting and almost four times more likely to report revenue growth, according to the survey. This is not a compliance exercise that gets reviewed quarterly. It is an operating system that assigns ownership of AI outcomes to specific leaders, creates measurement standards that make AI provable under scrutiny, and runs continuously alongside the AI it now governs.

2 **Close the C-suite alignment gap before it closes you**

CIOs/CTOs, COOs and CFOs see the organization differently. Until leadership shares a common definition of AI readiness, accountability and risk management, workforce investment will underperform and agentic deployments will scale without the controls to contain them.

3 **Measure what is working — and exit what is not**

Only 22% of operations leaders are working with a fully developed and implemented AI strategy. The organizations pulling ahead are not scaling more pilots — they are scaling fewer, with better measurement and clearer exit criteria. Depth creates the outcomes that justify the next investment.

ARE YOU IN THE

AI proof gap?

Five questions every executive must answer:

- 1 Do your leaders share a common definition of AI success, risk and accountability?
- 2 Can you consistently measure ROI across your AI initiatives and identify which ones should scale or stop?
- 3 Have you defined where AI should act autonomously, where human oversight is required and who is responsible for outcomes?
- 4 Could you produce auditable evidence of how your AI systems make decisions today?
- 5 If an AI system failed tomorrow, do you have a tested response plan and the ability to trace what went wrong?



If you answer “no” to any of these questions,
you are in the AI proof gap.

The organizations closing the proof gap are not waiting for better technology or external pressure to force the issue. They are building the governance that makes AI perform successfully today. The gap between those organizations and the ones still planning is already measurable. It will not close on its own.

Visit gt.com/ai to learn more
about how Grant Thornton
helps organizations build.

METHODOLOGY

Between February 23 and March 18, 2026, Grant Thornton conducted a survey of 950 business leaders across 10 industries and private equity. Respondents were drawn from senior leadership, including C-suite executives (CEOs, CFOs, COOs, CIOs/CTOs) and leaders reporting directly to the C-suite.

Respondents came from asset management (N=100), banking (N=50), construction/real estate (N=100), energy (N=100), insurance (N=100), manufacturing (N=100), media and entertainment (N=100), private equity fund leadership (N=100), services (N=100) and technology and telecommunications (N=100).

Functional representation came from operations (390), finance (313), IT (234) and CEO/managing partner (13).



Grant Thornton

GT.COM

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.

© 2026 Grant Thornton Advisors LLC | All rights reserved | U.S. member firm of Grant Thornton International Ltd.