Deloitte.

Together makes progress

NO VULNERABILITIES?

How you respond to frontier AI may determine whether it overwhelms your security program or strengthens it



Could we be heading toward zero vulnerabilities in the age of frontier AI?

It sounds absurd, especially when considered against the backdrop of today's cybersecurity reality, in which identified vulnerabilities can often accumulate faster than human teams can remediate them.

Add frontier models that can identify and act on vulnerabilities across codebases at a scale and speed most humans can't match, and if remediation stays human-paced, the resulting backlog can quickly overwhelm the team responsible for clearing it. Frontier AI models can apply sophisticated AI capabilities across software systems, identify subtle logic flaws that may have otherwise evaded teams of human reviewers, and help develop exploit techniques faster than defenders may be able to respond.

The exposure isn't limited to IT. Operational technology (OT) can carry the same vulnerabilities with an even smaller margin for error. In some scenarios, a single skilled attacker with access to the right model can now do what once required a large team and deep institutional knowledge.

So how could anyone even consider the possibility of zero vulnerabilities in this environment?

Because the same capability that promises to overwhelm cybersecurity teams today is the one that could flip this trend on its head. Cybersecurity teams may have access to the same frontier AI models as the cyber attackers—and can deploy them across the full security lifecycle, from accelerating remediation to reducing the likelihood that vulnerable code reaches production in the first place. Just as important, these capabilities can help remediate what's already in place faster than traditional approaches.

Of course, that may not happen without a shift in approach. But it is a possibility—one that's intriguing enough to seriously examine. When attention is centered on the immediate demands of scanning, patching, and triaging, incorporating AI models to help prioritize vulnerabilities and focus remediation efforts where they matter most can be challenging. That work is urgent and necessary—but it may not be sufficient given the scale of change that frontier AI may introduce. Frontier AI capabilities may soon be standard equipment for both sides—attackers and defenders. The bigger challenge is to build the organizational speed to help act on what these models find. Organizations that don't use this lead time to build frontier AI capabilities into their development and remediation processes may fall behind on both fronts.

When organizations look back on this moment in a few years, they're likely to conclude that those that thrived when confronted by the challenge of frontier AI were the ones that proactively used these capabilities to embed security into how they build and operate, rather than bolting it on after the fact. This technology advancement may provide a similar opportunity to what was seen during the era of large-scale cloud migration. Many organizations may have missed their first, best opening to embed security as they moved to take advantage of the cloud and ended up scrambling to retrofit their security operations after the fact. Security was not necessarily ignored; however, in the push to move quickly, capture cost savings, and modernize the stack, it may not always have been integrated early enough.

Comprehensive, infrastructure-level security was too often layered in once the migration was complete. The lesson is straightforward: embedding security early can help manage the downstream cost and complexity of retrofitting it later.

Frontier AI may be giving companies the opportunity to get the foundations right while everything is in motion. Here's how to think about moving ahead.





Why not just wait?

Today, many organizations can already access the most capable frontier AI models, and that access will only expand. The real challenge is finding ways to successfully deploy them at scale. As these capabilities commoditize, even less sophisticated threat actors—those without the knowledge to mount sophisticated attacks on their own—may gain the same access. It is possible that significant investment by nation-states in their own AI capabilities could further accelerate that timeline.

Companies that don't use this time to change how they build and defend may find themselves on the wrong side of a capability gap that may be very difficult to close. Waiting may also compound the cost of eventual action.

Starting now may give organizations more opportunity to build relevant expertise over time, rather than having to develop it under greater pressure later. The knowledge that matters may be changing, too. Finding vulnerabilities may be becoming a commodity skill. The harder, more valuable skill may be knowing which vulnerabilities matter most to the business—and what to do about them.

There may still be time to benefit from early embedding. Those who wait should expect to pay more, inherit more technical debt, and perform catch-up work under mounting pressure.

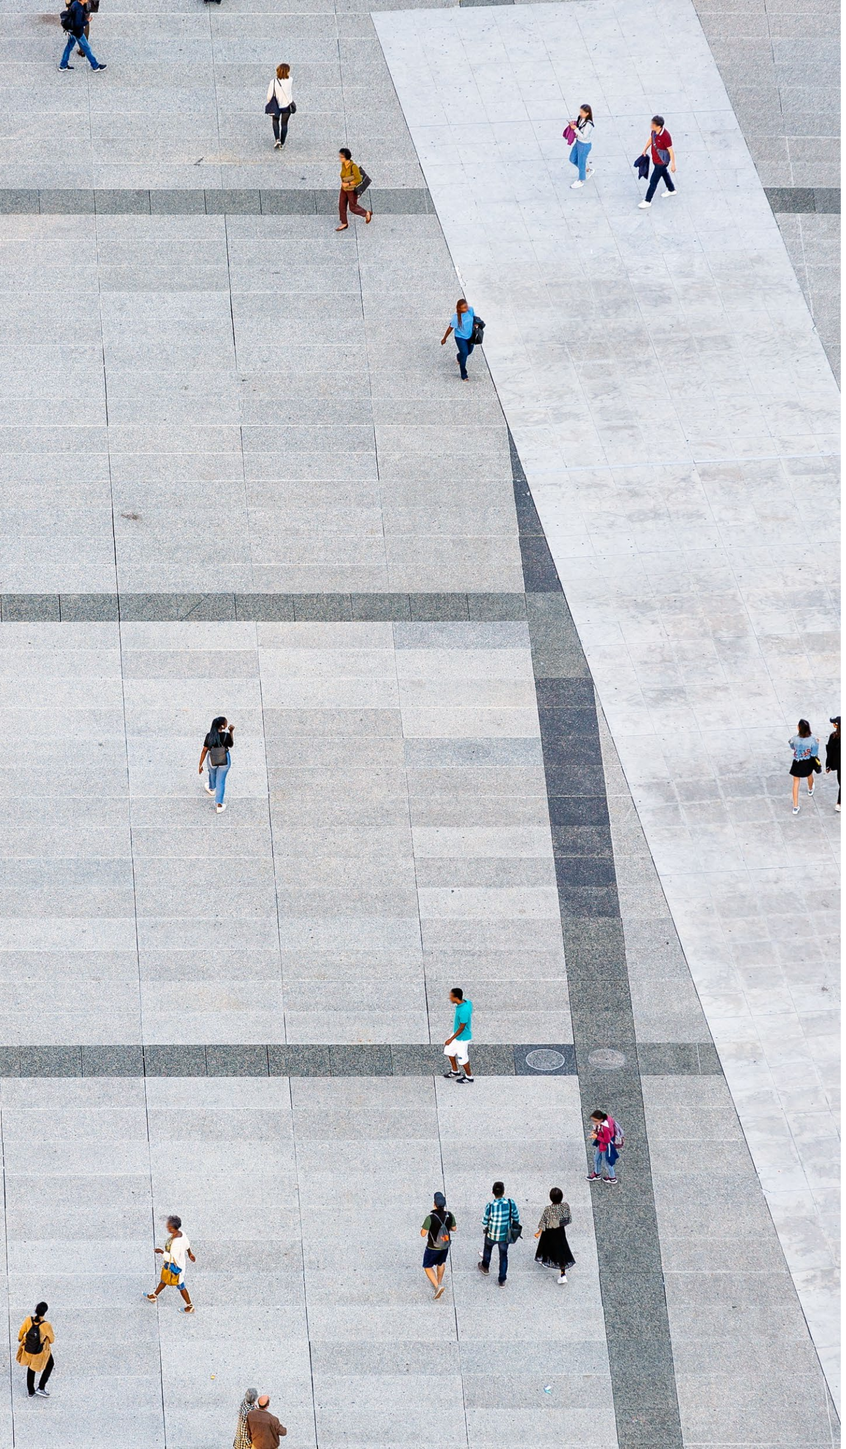

The remediation trap

Frontier AI models may already be forcing many organizations to do the right thing for current problems. They may be scanning faster, triaging more systematically, and working to reduce time to remediation. All important work—but it treats the symptoms, not the underlying causes.

The numbers can be daunting. In a reported deployment, security teams using a frontier AI model achieved vulnerability-detection rates more than 10 times higher than their previous human-led and automated baselines[1]. If tools for finding vulnerabilities become substantially faster than the processes used to fix them, patching faster could become a treadmill rather than a finish line.

The push to remediate faster may carry a less-obvious risk. Patches applied too fast, without adequate testing, can break production systems. The cure can create new vulnerabilities and other potential operational challenges. Frontier AI can help here, accelerating remediation while maintaining the quality controls that high-velocity patching can compromise. But even with AI assistance the pressure to close tickets can outpace the judgment required to close them *safely*. As frontier AI capabilities become more widely available, the deciding factor won't be who has access to the best models—it may be who can act quickly on what those models find. Speed should be embedded in security upstream to help avoid this trap.

Legacy debt offers one lens for understanding the challenge, particularly where accumulated workarounds make systems difficult to rearchitect and could be contributing to the backlog. It's possible that many of the vulnerabilities frontier AI is uncovering aren't new. They may have been there for years, quietly waiting. Addressing them through patching is like painting over rotten wood—it changes the surface without changing the structure. Deloitte's 2026 Global Technology Leadership Study estimates that technical debt accounts for 21–40% of IT spending, and notes that substantial value remains trapped in existing technology investments. The question is whether certain legacy systems can be worth continuing to defend, or whether the resources devoted to keeping them patched may be better spent on replacement.

TECHNICAL DEBT WITHIN IT SPENDING:

21–40%

The organizations escaping the remediation trap may not be patching less. But they may be focusing on a different question as they go: "How do we minimize the number of vulnerabilities that reach production in the first place?" This is the shift from reactive to proactive—from patching code to helping to secure the process that produces it.

1 Anthropic, "Project Glasswing: 10,000+ Critical Vulns Found in a Month," *KubeSimplify Diaries*, May 25, 2026.

Making the business case

The conversation most CISOs should be having today is about capacity and prioritization, not just budget. Adding headcount or procurement line items may not solve the structural problem if the underlying workflows, incentives, and governance structures can remain unchanged. What is the organization willing to change about how it operates to make security in a frontier AI world a priority in everything it builds? The answer may require cross-functional buy-in from engineering, product, operations, and—importantly—from the C-suite and board. The CISO can raise the issue and frame the risk. But they may not change the organizational model and institutional culture on their own.

AVERAGE BREACH COST IN THE US:

$11.5 million

Framing this for a CFO or board requires translating warning into cost. Globally, the average cost of a data breach climbed to nearly $5 million in 2026, up 12% year-over-year. In the US, the average breach cost reached a record $11.5 million, driven by regulatory penalties and slower detection times.[2]

Organizations without access controls for AI systems accounted for 92% of AI-related breaches.[3] Organizations that embed AI into their defenses may face fewer of these blind spots. Those that don't may remain exposed to both rising threats and the costs that come with them. In many cases, proactive investment in embedded security can cost less than reactive remediation following a breach.

Boards and senior leaders need to recognize that inaction is a choice with a cost. Every quarter spent without a clear frontier AI security strategy can compound the risk. Agreeing on a frontier AI strategy may not be enough in this context. The harder work may be driving strategic priorities into the business—embedding frontier AI-ready security into workflows and operations. Executing the translation of strategy into practice can be where most organizations typically fall short facing this type of challenge.



2 IBM, Cost of a Data Breach Report 2026 (July 29, 2026), https://www.ibm.com/reports/data-breach.

3 "IBM's 2026 Data Breach Report: 92% of AI Incidents Had No Access Controls," Cybersecurity Insiders, August 3, 2026, https://www.cybersecurity-insiders.com/ibm-2026-data-breach-ai-access-controls/.

What embedding really means

Embedding frontier AI into operations may sound smart enough. But what does it mean, really? In practice, it can mean integrating AI-powered security capabilities into the workflows that help create, deploy, maintain and patch software, rather than running them in parallel or applying them after the fact. Embedded security happens at the point of creation, which is the only point where it can change the economics of the vulnerability problem. So does AI-assisted maintenance of what's already in production.

What embedding really means



ASSESS

Assessment can require understanding your current governance posture, regulatory obligations, existing AI vendor relationships, and organizational readiness to absorb change. This may not be a CISO-only exercise. Business leaders should be identified, engaged, and held accountable for security outcomes in their domains. An assessment may identify a gap between the adoption of AI tools and the organization's ability to maintain visibility and apply consistent controls. And while most organizations may know their technology stack reasonably well, fewer may have a clear picture of the security assumptions embedded in that stack—or the places where those assumptions no longer hold. Addressing this vulnerability can require CISOs and business leaders to work together to help ensure accountability and execution.

AI should be embedded in the assessment work itself, used to find and prioritize vulnerabilities. For example, it can be used to determine business impacts, identify dependencies based on resiliency objectives, and weigh exploitability quickly and effectively compared to manual triage approaches that may not be able to keep pace with the volume of findings these environments generate today.



SECURE

This is where AI may get truly embedded into how the organization operates and protects itself. A secure Software Development Life Cycle (SDLC) helps to ensure security is integrated into development from the first line of code. Continuous threat exposure management may replace point-in-time assessments with an ongoing picture of where the organization is vulnerable. Identity and access hygiene limits the blast radius of any individual compromise. Data protection disciplines help ensure that what AI models can access, store, and transmit is governed deliberately. Integration with operational resilience planning can turn security from a separate workstream into a core part of how the organization prepares for and responds to disruption.



DETECT AND DEFEND

Once systems are built and deployed, the work of protecting them should be continuous. In this phase, AI augments human analysts across the full operational cycle—triaging alerts, identifying patterns across disparate data sources, and flagging the signals that require human judgment.

But detection is only the start. AI can accelerate threat assessment by correlating indicators of compromise perhaps even faster than any human team, and supports response by recommending (or even initiating) containment actions before damage can spread. Organizations that have done the embedding work in earlier phases may find that this operational layer performs significantly better, because the AI has context, the integrations are in place, and the humans know how to work with it.

The right entry point in this framework can vary by organization. A company with mature governance and relatively modern infrastructure might begin with the secure dimension—embedding AI directly into engineering workflows. A company with significant legacy debt and regulatory complexity might need to start with a thorough assessment before any other investment makes sense. Either way, building security in and maintaining what already exists may not be competing priorities—they can be the same work, done with the same tools, at different phases of the lifecycle.

Is zero vulnerabilities really achievable?

That’s not the point.

You don’t have to believe that a zero-vulnerability enterprise is achievable to benefit enormously from the process of attempting to get there. It may not be possible—but the goal here is directional, not literal. Even short of perfection, limiting business harm and keeping operations running when prevention efforts fall short are goals that can be worth pursuing. You can still land with cleaner code, stronger governance, and a more defensible posture. Just as important, you’ll have an organization that can absorb a hit and keep moving without breaking. The downside of aiming too high is modest. But the downside of not trying (or *not trying* fast enough) is severe.

In a landscape fraught with frontier AI, making the organization significantly harder to breach isn’t a prize for missing an impossible target—it may be the target.

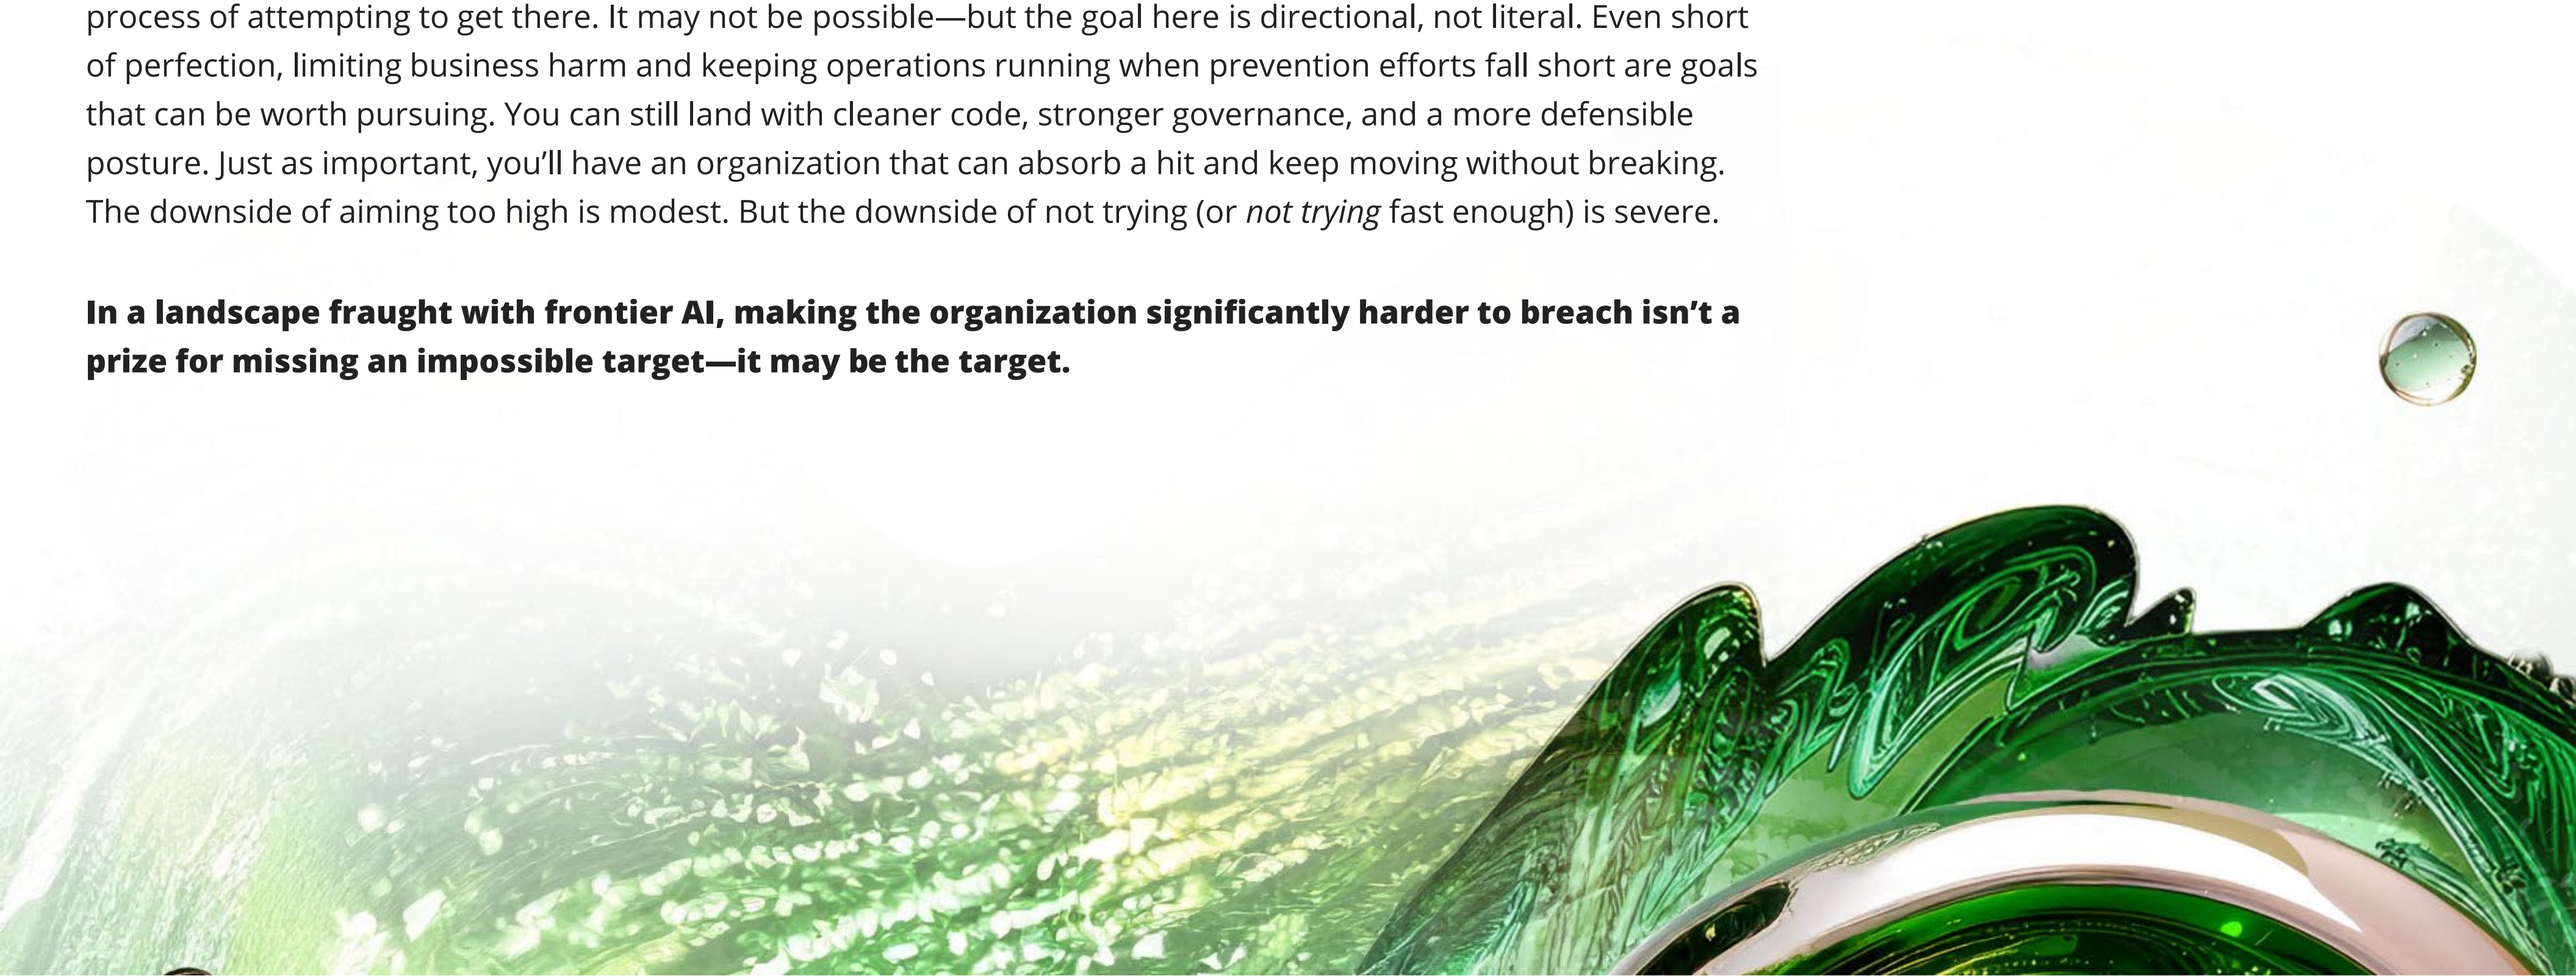

CONTACTS



Ayan Chatterjee
Deloitte Global Cyber Leader
Principal, Deloitte Consulting LLP
+1 714 436 7157
globalcyberleader@deloitte.com



Adnan Amjad
US Cyber Leader
Partner, Deloitte & Touche LLP
+1 713 982 4825
aamjad@deloitte.com



Nicola Esposito
EMEA Cyber Leader
Partner, Deloitte Spain
+34 918232431
niesposito@deloitte.es



Yuichiro Kirihara
Asia Pacific Cyber Leader Partner,
Deloitte Japan
+81 803 3672805
ykirihara@tohmatsu.co.jp



Daphne Lucas
Canada Cyber Leader
Partner, Deloitte Canada
+1 403 267 1737
dalucas@deloitte.ca



Paula Alvarez
S-LATAM Cyber Leader
Partner, Deloitte Mexico
+52 55 5080 6558 ext: 6558
palvarez@deloittemx.com



Jose Pela Neto
Brazil Cyber Leader
Partner, Deloitte Brazil
+55 11 5186 6396
jpela@deloitte.com



Criss Bradbury
Cyber Asset Leader
Principal, Deloitte & Touche LLP
+1 303 305 3218
cbradbury@deloitte.com



Niels van de Vorle
Cyber Market Growth Leader
Partner, Deloitte Netherlands
+31 882 882 186
nvandevorle@deloitte.nl

Deloitte.

Together makes progress

Disclaimer

Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited (DTTL), its global network of member firms, and their related entities (collectively, the "Deloitte organization"). DTTL (also referred to as "Deloitte Global") and each of its member firms and related entities are legally separate and independent entities, which cannot obligate or bind each other in respect of third parties. DTTL and each DTTL member firm and related entity is liable only for its own acts and omissions, and not those of each other. DTTL does not provide services to clients. Please see www.deloitte.com/about to learn more.

Deloitte provides leading professional services to nearly 90% of the Fortune Global 500® and thousands of private companies. Our people deliver measurable and lasting results that help reinforce public trust in capital markets and enable clients to transform and thrive. Building on its 180-year history, Deloitte spans more than 150 countries and territories. Learn how Deloitte's approximately 460,000 people worldwide make an impact that matters at www.deloitte.com.

This communication contains general information only, and none of Deloitte Touche Tohmatsu Limited (DTTL), its global network of member firms or their related entities (collectively, the "Deloitte organization") is, by means of this communication, rendering professional advice or services. Before making any decision or taking any action that may affect your finances or your business, you should consult a qualified professional adviser.

No representations, warranties or undertakings (express or implied) are given as to the accuracy or completeness of the information in this communication, and none of DTTL, its member firms, related entities, employees or agents shall be liable or responsible for any loss or damage whatsoever arising directly or indirectly in connection with any person relying on this communication. DTTL and each of its member firms, and their related entities, are legally separate and independent entities.

© 2026. For information, contact Deloitte Global.