Produced by the AI Governance Pillar at AI Singapore,
now part of the NUS Artificial Intelligence Institute

AI Governance Roundtable #5:
AI Adoption Among Non-tech Corporates

How can governance support AI adoption among non-tech corporates? How can industry and corporations navigate the proliferation of AI governance across multiple countries and regions?

This is the fifth of a series of roundtables convened by AI Singapore for representatives from industry, government, and academia to discuss responsible AI. Such discussions are typically too narrow and too broad. Too narrow in that a few voices dominate the discussion – notably those in the United States and Europe, with China sometimes included. Too broad, in that the discussion is often limited to generalities and principles. This project aims to address both aspects of this problem, involving a wider set of stakeholders — in particular those from Southeast Asia — in more focused discussions of specific challenges in the application of Responsible AI to particular questions.

Rapporteur: Hakim Norhashim, AI Singapore
Location: Deloitte Singapore

Executive Summary

In November 2024, AI Singapore convened a roundtable at Deloitte Singapore to examine how governance can better support AI adoption among non-tech corporates, particularly in light of the growing regulatory complexity across jurisdictions. As AI systems become increasingly embedded in business operations, organisations face the challenge of adopting AI responsibly while navigating a fragmented global governance landscape. **This is particularly pertinent for**

non-tech corporates, who often play no role in building these systems, yet are expected to interpret complex regulations, manage risks, and take responsibility for the outcomes.

Participants in this roundtable also emphasised the crucial role that Small and Medium Enterprises (SMEs) play in the local and regional South-East Asian economy and highlighted how SMEs can be further empowered and, in some instances, lead AI innovation.

The roundtable brought together stakeholders from industry, regulatory bodies, and academia to explore the obstacles that hinder responsible adoption and to surface operational insights that could inform future policy, operations, and cross-sector collaboration. This report thus reflects the key insights, structured as *AI Governance Gaps*, that have emerged throughout the discussion, providing a synthesis of the gaps raised and the proposed solutions.

Discussion Themes

AI Governance Gap 1: AI Governance Measures Fail to Inform Real-World Operations
Governance frameworks often remain too abstract to effectively guide AI workflows, leaving non-tech corporates uncertain about how to act, which results in either shallow compliance or inaction. This challenge is compounded for organisations operating across jurisdictions, where inconsistent regulatory expectations increase interpretive burdens. What's needed is a translation workflow that interfaces with existing AI operations, linking high-level principles to day-to-day decisions and reconciling external complexity with internal clarity.

AI Governance Gap 2: SMEs Are Expected to Adopt AI in Ecosystems Not Designed to Support Them
Across Southeast Asia, small and medium-sized enterprises (SMEs) constitute 99% of all firms and contribute significantly to employment and GDP across the region.[1] While SMEs are broadly encouraged to adopt AI responsibly, they face ecosystems that assume resources they often lack, such as legal counsel, in-house AI expertise, or robust digital infrastructure. Without structural support, many are forced to choose between incurring excessive risk or opting out of adoption altogether. Given the pervasive nature of AI disruption and the outsized role SMEs play in sustaining regional economies, supporting SMEs is no longer optional; it is a strategic imperative.

AI Governance Gap 3: Domain Experts Are Excluded from AI Operations
AI systems often remain overly technical and are built without clear pathways for domain experts to intervene, resulting in pipelines that are technically sound but contextually lacking. SMEs often face an all-or-nothing AI proposition, where they must adapt their processes to fit off-the-shelf systems with limited customisation. However, if AI governance is to be truly contextual, it cannot be top-down alone; it must include input from those closest to operational impact. AI governance thus has a critical role in embedding human judgment from various

[1]ERIA/OECD (2024), SME Policy Index: ASEAN 2024 – Enabling Sustainable Growth and Digitalisation, Jakarta: ERIA/Paris: OECD.

domains, including operations, product, communications, legal, and compliance, into AI workflows and at the point where it matters most: in real-time, operational decision-making.

AI Governance Gap 4: AI Governance Is Often Framed as a Constraint
Governance remains framed as compliance-driven, alienating non-technical teams and discouraging innovation. This challenge is compounded by the fragmentation of AI governance across jurisdictions, where divergent narratives, precautionary in the EU, market-first in the US, and state-centric in China, reinforce the perception that governance is inconsistent and externally imposed. For non-tech corporates, particularly those operating across markets, this ambiguity erodes confidence and reinforces the view that governance is something to be avoided, rather than engaged with. Reframing governance as a strategic enabler, supported by an internal translation workflow that reconciles external expectations with internal clarity, can transform the paradigm of governance from a *constraint* that hinders AI adoption to a *catalyst* that empowers collective AI innovation throughout the organisation.

Taken together, these four gaps reveal a critical insight: the prevailing paradigm, in which AI governance is regarded as merely a matter of compliance or a hindrance to innovation, needs to be reimagined. Governance needs to evolve into a driving force for AI innovation, rooted in the inclusion of non-tech corporates in conversations across the whole AI pipeline.

The discussion extended beyond identifying gaps, as participants actively suggested solutions. A core theme throughout was the need to move away from externally imposed AI governance toward governance that is developed with and through organisations.

Recommendations at a Glance

1. **Organisations Adopting AI: Incorporate a *Translation Workflow* to Operationalise AI Governance – Bridging Regulations, Policies and AI-Enabled Operations**
   All organisations adopting AI should establish dedicated translation workflows to bridge high-level AI governance frameworks with day-to-day operations that are impacted by AI. This is not a new management tier or compliance bureaucracy. Rather, it is a functional architecture: a dynamic, cross-functional workflow that ensures governance is embedded at the point of real-world impact, easing interpretation, aligning regulatory intent with practical tools, and connecting internal teams to external support.

   Effective governance depends on striking the right level of abstraction by balancing top-down principles with bottom-up insights on deployment. For less AI-mature organisations, particularly SMEs, this balance can be achieved through co-developing contextualised governance scaffolding with experienced ecosystem partners. Over time, this translation workflow should feed into a learning pipeline that captures, shares, and refines implementation insights, contributing to a collective repository of successful governance use cases and supporting both internal alignment and broader ecosystem learning.

An important aspect of operationalising AI governance discussed at the roundtable was the role of AI governance tools. Essentially, tools are how principles become practices. They provide the procedural and technical scaffolding through which abstract governance ideals, such as fairness, transparency, and accountability, are made operational. The translation workflow plays a key role in this process by serving as a decision-making juncture: it identifies where tools are needed, selects which are appropriate for the organisation's context, and integrates them into AI operations. From model and system cards to documentation aids, output review mechanisms, prompt control safeguards, and audit processes, these tools render governance interpretable and actionable, especially for non-technical teams. When shared by experienced ecosystem players or co-created within regulated sectors, tools translate principles into safeguards, lowering the barrier to adoption and enabling teams, especially SMEs, to implement responsible AI with clarity and confidence.

2. **Regulators, Banks, Tech Companies, and Industry Associations: Design the AI Ecosystem Around SME Enablement**
   Support for SMEs must be intentional and by design. A healthy AI governance ecosystem distributes responsibility across multiple actors. Regulators can provide context-aware guidance tailored to diverse use cases. Banks can play a role in enabling trusted onboarding and driving sector-wide learning. Tech firms can contribute by opening access to tools, infrastructure, and governance scaffolding. Industry associations can help localise standards and foster sector-specific communities of practice. Collectively, the burden on SMEs can be reduced, creating the conditions for meaningful and responsible AI adoption across the economy.

3. **SMEs: Lead real-world pilots for AI applications and governance**
   SMEs are well-positioned to lead the experimentation that responsible AI adoption demands. But to do so meaningfully, governance systems must be intentionally designed around their constraints and strengths, particularly their agility, contextual awareness, and sector-specific expertise. Positioning SMEs as first-movers in use case development creates space for grounded, real-world testing of both AI capabilities and governance mechanisms. This must be supported by ecosystem players who can anchor communities of practice and share both tech and domain-specific learning. Participants also called for the creation of a shared repository of successful and unsuccessful AI use cases—a collective resource that can reduce risk and accelerate learning. With structured pilot testing and shared infrastructure, perhaps SMEs will not only represent 99% of firms in the broader Southeast Asian economy, but also 99% of their AI use cases and applications.

4. **Organisations Adopting AI: Embed Organisational Domain Knowledge Into AI Workflows**

Similar to the first recommendation on translation workflows, this recommendation addresses the need to make governance tangible within day-to-day AI use. While translation workflows focus on interpreting external governance frameworks and aligning them with internal operations, this recommendation focuses on embedding organisational knowledge, including legal, operational, and product-specific knowledge, directly into AI workflows.

AI pipelines must be examined and restructured to incorporate domain expertise at key points where human judgment can meaningfully influence, adjust, or, if necessary, constrain outputs. Even when AI systems are fundamentally opaque, embedding human judgment throughout the AI workflow, for example, in the form of AI governance tools, which form the interface between domain experts and AI workflows, enables organisations to assess the feasibility of AI deployment, scrutinise outputs, and make risk-informed decisions grounded in context. These intervention points are critical for maintaining contextual relevance and ensuring real-time accountability. Other technical methods, such as fine-tuning, prompt controls, and reinforcement learning with human feedback (RLHF), can also make interventions to modify AI outputs feasible in practice, even if the overall model behaviour remains partially opaque. This enables human expert intervention to be the means through which organisations can effectively manage the limitations of AI's black box nature.

5. **Governments, Regulators and AI Governance Experts: Frame AI Governance Not as a Constraint, but as a Catalyst of Collective AI Innovation**
   AI governance should not be perceived as a constraint, but rather as a "catalyst" that empowers collective AI innovation across the organisation. Framing matters: how regulators and organisations talk about governance shapes who participates. When governance is positioned as a support structure, it becomes a bridge between technical teams, business units, legal, communications, and leadership.

Ultimately, the answer to the question of how governance can support AI adoption begins with the inclusion of non-tech corporates in AI conversations. Non-tech corporates, particularly in SMEs, are not just end users of AI; they are domain experts, accountable decision-makers, and stewards of real-world impact. This report shows that governance needs to transform from compliance into a collaborative organisational capability, one that shapes AI adoption around what is meaningful and valuable. While technical stakeholders will continue to advance the frontier of AI research, non-tech corporates have a unique opportunity to lead where AI truly matters: in its impact on people's lives and its alignment with the values we live by. The question then is no longer whether governance will support AI adoption; it is how deliberately and collectively we are willing to shape our shared future with the empowerment of AI.

Contents

About the Roundtable: Supporting AI Adoption Among Non-tech Corporates Amid Evolving Global AI Governance

In November 2024, AI Singapore convened a roundtable at Deloitte Singapore to explore two interrelated questions: *How can governance support AI adoption among non-tech corporates?* and *How can industry and corporations navigate the proliferation of AI governance across multiple countries and regions?* These questions were framed in response to the growing demand for responsible AI adoption across sectors, particularly among organisations that may not be AI-native but are increasingly developing, deploying, or relying on AI technologies. At the same time, businesses must navigate a rapidly evolving patchwork of AI-related regulations and principles across jurisdictions, creating uncertainty around risk, compliance, and best practices. This is especially true across Southeast Asia, where many non-tech corporates, especially SMEs, face similar structural and regulatory challenges amid growing pressure to digitalise.

The two-hour discussion brought together a small group of representatives from industry, regulatory bodies, and academia. The session was facilitated by Professor Simon Chesterman and conducted under the Chatham House Rule. Rather than formal presentations, participants were invited to share brief remarks outlining their understanding of the questions at hand, any tentative answers, and what kinds of regulatory clarity or academic insight would support effective AI adoption and governance. This report reflects the key themes and insights that have emerged, providing a synthesis of the challenges raised, the solutions proposed, and the perspectives exchanged.

Discussion Themes: Addressing Governance Gaps to Support AI Adoption

The discussions at the Deloitte Roundtable on AI Governance centred around four key governance gaps that reflect the institutional, operational, and cultural barriers faced by non-tech corporates in adopting AI. These themes surfaced through rich exchanges among participants across regulatory, industry, and academic sectors, each contributing insights grounded in practical operational realities.

While the roundtable was anchored in the Singapore corporate context, many of the challenges raised, such as translating principles into workflows, supporting SMEs, embedding domain expertise, and reframing governance, resonate across regions and organisational types. In particular, participants in this roundtable emphasised the crucial role that SMEs play in the economy, locally and across ASEAN, and highlighted how SMEs can be further empowered and, in some instances, lead AI innovation. Together, the themes highlight structural issues that hinder AI adoption, while also revealing actionable opportunities to reimagine governance as a structural catalyst that enables, rather than constrains, AI innovation.

The following sections present each of the four governance gaps in detail, accompanied by targeted recommendations drawn directly from the discussion.

AI Governance Gap 1: AI Governance Measures Fail to Inform Real-World Operations

AI governance frameworks, ranging from hard regional regulations like the EU AI Act to soft company policies, often focus on broad, high-level principles. While principles are necessary to anchor AI operations in desirable overarching values as well as to set ethical boundaries, participants highlighted significant obstacles, most notably the inability to translate such abstract guidance into operationalisable governance interventions. This is especially the case for organisations navigating multiple regulatory jurisdictions, such as the EU AI Act, Canada's ADA, or national standards in China and the US, where inconsistent expectations increase the difficulty of implementation and exacerbate interpretive burdens. As one participant observed, *"A lot of these governments are going down a very blanket approach... It's going to be very hard for companies to understand what they need to comply with unless there is interoperability between these different jurisdictions."* This gap is particularly pronounced among non-tech corporations and SMEs, which, though well-versed in compliance within their own sectors, are navigating unfamiliar ground when it comes to AI. Domain expertise and established processes do not necessarily translate into fluency with the emerging, fast-moving concerns specific to AI governance. As one participant noted, "It's difficult when there are unknown unknowns… and this area [AI governance] is changing so quickly… it's hard just to keep up, and for SMEs in particular, it's understandable that there's a lot of weariness, a lot of concern."

This failure to bridge principles with operational reality is further complicated by the fact that governance processes are often described as “very complicated, expensive, and difficult,” which pushes SMEs toward either superficial compliance or complete inaction. Many struggle to translate abstract principles into concrete workflows because the guidance itself is vague and lacks procedural precision. A recurring theme during the discussion was the need to “move beyond generalities,” alongside a parallel call for regulatory clarity. Governance, participants noted, must rest on concrete guidance rather than leave parties to infer intent from ambiguous expectations. One specific suggestion was for governance to be structured explicitly around clear, practical use-cases, with one participant asserting that if SMEs can see clearly that “the regulator’s concerns are with regards to use-cases ABC,” then they can confidently “do XYZ.” Without this level of specificity, bridging the gap between guidelines and AI operations places an undue burden of interpretation on stakeholders who may not be as familiar with the technology. This further risks relegating AI governance to a theoretical exercise where outcomes are not only uncertain but may also lack standardisation, rather than serving as a means to effectively mitigate risk.

One participant succinctly framed these risks as liabilities, noting that SMEs need assurance that their actions are safe or at least compliant; thus, if something goes wrong, they can demonstrate that they have fulfilled their due diligence through adhering to available guidelines. On the other hand, another participant emphasised that the current state of ambiguity directly leads to "uninformed use or misapplication" because SMEs "simply lack the confidence and the tools to evaluate if [AI] is right for them." Indeed, what’s clear is the need to “move beyond generalities”, both in how technically proficient actors in the AI ecosystem provide meaningful support to others, including SMEs, and in the expectations regulators place on downstream deployers of AI. Greater clarity on both fronts will better equip in-house experts to translate external expectations into actionable guidance for internal teams.

Recommendation 1: Incorporate a Translation Workflow to Operationalise AI Governance, Bridging Regulations, Policies and AI-Enabled Operations

For Organisations Adopting AI

To bridge this gap, participants called for operationalising existing governance principles by incorporating a critical *translation workflow for governance*. This workflow will effectively transform overarching principles into concrete, actionable procedures suitable for implementation within the organisation. As one participant succinctly noted, "Principles aren't enough. Principles will get you only so far." The role of this translation workflow includes navigating the complexity of multiple regulatory regimes across jurisdictions. It is important to note that regulations must be interpreted specifically for each AI application. Thus, there is no overarching solution that works across all jurisdictions for all use cases. This is precisely where the translation workflow becomes essential: it provides the interpretive scaffolding needed to tailor governance to real-world operational contexts while remaining compliant with fast-evolving regulatory environments.

This translation workflow must address two core asymmetries: first, the disproportionate interpretive burden placed on SMEs and non-technical stakeholders in larger organisations who lack the fluency to extract operational meaning from vague guidelines; second, the disconnect between regulatory expectations and the guidance currently available to downstream deployers of AI. Clarity in both areas is necessary not only for compliance but to enable meaningful participation in AI governance. Participants emphasised that AI governance should be based on clear guidance, rather than relying on parties to deduce intent from ambiguous expectations.

The key to getting the translation right is achieving the appropriate level of abstraction in operationalisation. Guidance must be concrete enough to inform day-to-day decisions yet flexible enough to accommodate varied sectors, diverse use cases, and the inherent variability of business-as-usual (BAU) operations. After all, AI governance should not become a barrier to how organisations function on the ground. To strike this balance, participants emphasised the need for guidance to be informed by operational realities, moving towards contextualised, sector-relevant scaffolding that departs from static templates or checklist-style governance. The goal is to build tools that make governance adaptable to business context and operational workflows, not the other way around.

One participant suggested establishing jumpstart programs to help SMEs build out successful implementations. At the same time, another emphasised the value of larger organisations sharing best practices with smaller enterprises, as this could "automatically help to raise the level of AI literacy across the economy." Others described how peer-led use case references, such as in the accounting sector, made adoption more feasible, especially when practitioners could see how others successfully implemented AI. These supports are essential for equipping in-house teams to interpret external expectations and embed them into governance processes that make sense within their specific domain. Crucially, the development of such contextualised scaffolding should not fall solely to individual organisations but rather be co-developed by regulators, ecosystem partners, and technically proficient actors who can bridge the gap between policy and operations.

An important aspect of operationalising AI governance discussed at the roundtable was the role of AI governance tools. Essentially, tools are how principles become practices. As one participant observed, "Principles aren't enough. Principles will only get you so far." What organisations need are tools and mechanisms to act on those principles. Tools provide the procedural and technical scaffolding through which abstract governance ideals, such as fairness, transparency and accountability, are made operational.

The translation workflow plays a key role in this process by serving as a decision-making juncture. It identifies where tools are needed, selects those that are appropriate for the organisation's context, and integrates them into AI operations. As another participant noted, "In the absence of any rules, trust is enabled through knowing that the process will be the same each time," underscoring how consistency and tooling can reinforce confidence across teams.

From model and system cards to documentation aids, output review mechanisms, prompt control safeguards and audit processes, these tools render governance interpretable and actionable, especially for non-technical teams. As one participant put it, "It really should answer questions A, B, C, D, E, F," highlighting how well-designed documentation tools can clarify expectations and reduce uncertainty.

When shared by experienced ecosystem players or co-created within regulated sectors, tools help translate principles into safeguards. As one contributor from an established stakeholder group put it, "We will let you know when it works… and when it doesn't." These shared resources lower the barrier to adoption and empower teams, particularly SMEs, to implement responsible AI with clarity and confidence.

Importantly, the translation workflow should not be viewed as a narrow technical or compliance function; rather, it is a structural necessity that embeds governance considerations from the outset, rather than treating AI governance as a downstream, reactive, or purely compliance or regulatory concern. Without it, organisations risk acting without confidence, falling into either inaction or superficial compliance. When well-designed, this workflow becomes the bridge that enables non-tech actors to adopt AI responsibly.

In summary, participants strongly advocated embedding governance into operational practice. Only through precise operationalisation, supported by regulators as well as technically proficient ecosystem actors, can AI governance fulfil its intended role: to empower informed, effective, responsible adoption across the full spectrum of organisations.

AI Governance Gap: AI Governance Measures Fail to Inform Real-World Operations

Recommendation: Organisations Adopting AI — Incorporate a Translation Workflow to Operationalise AI Governance, Bridging Regulations, Policies and AI-Enabled Operations

1. **Incorporate a Translation Workflow to Bridge AI Governance Concerns and AI-Enabled Operations**

   Establish a dedicated function within your organisation that translates high-level AI Governance frameworks into operational interventions suited for BAU operations that are impacted by AI. This workflow should address three core functions: easing the interpretive burden on your organisation, aligning regulatory intent with available tools, and serving as a bridge between internal teams and external support.

2. **Get the Level of Abstraction Right: Co-Develop Contextualised Governance Scaffolding With More Experienced Partners**

Pinning the right level of abstraction is critical; without it, governance remains vague. Getting it right requires a balance between top-down principles and bottom-up insight from how AI is actually deployed in practice. When governance reflects both, it becomes contextualised, therefore actionable. But for organisations or sectors just beginning their AI journey, especially SMEs, the bottom-up view may not yet exist. In these cases, collaboration becomes essential. Working with regulators, ecosystem partners, and technically proficient actors allows these organisations to co-develop governance scaffolding grounded in operational reality, rather than abstract compliance. These scaffolds, in turn, should feed into a learning pipeline, capturing implementation feedback to continuously refine governance over time.

3. **Build a Learning Pipeline alongside a Repository of Successful AI Governance Use Cases**

   Governance should evolve through application. Develop mechanisms to capture, translate, and share lessons from implementation, both within organisations and across the sector. This learning pipeline should feed into a shared repository of successful AI governance use cases, allowing organisations to reference what works, what doesn't, and why. Internally, it enables internal teams, such as legal, risk, and product teams, to align over time. Externally, it builds ecosystem intelligence by making governance decisions, trade-offs, and design patterns visible. Rather than starting from scratch, organisations can learn from lived experience as well as contribute to a growing knowledge base that supports responsible AI adoption.

4. **Leverage AI Governance Tools**

   An important aspect of operationalising AI governance discussed at the roundtable was the role of AI governance tools. Essentially, tools are how principles become practices. They provide the procedural and technical scaffolding through which abstract governance ideals, such as fairness, transparency, and accountability, are made operational. The translation workflow plays a key role in this process by serving as a decision-making juncture: it identifies where tools are needed, selects which are appropriate for the organisation's context, and integrates them into AI operations. From model and system cards to documentation aids, output review mechanisms, prompt control safeguards, and audit processes, these tools render governance interpretable and actionable, especially for non-technical teams. When shared by experienced ecosystem players or co-created within regulated sectors, tools translate principles into safeguards, lowering the barrier to adoption and enabling teams, especially SMEs, to implement responsible AI with clarity and confidence.

AI Governance Gap 2: SMEs Are Expected to Adopt AI in Ecosystems Not Designed to Support Them

SMEs are often lauded as a core stakeholder in Singapore's economy, accounting for a majority of employment. While SMEs are broadly encouraged to adopt AI responsibly, participants noted that SMEs face a stark reality: the AI ecosystem they are expected to navigate was not built with them in mind. Most frameworks assume the presence of legal departments, compliance officers, and in-house AI expertise, resources that large firms may take for granted but which SMEs often lack, sometimes entirely. This gap in structural readiness was highlighted multiple times during the roundtable, with participants noting that SMEs are often excluded from design considerations, despite forming the vast majority of the economy. As one participant observed, "These matters matter to SMEs. We can't ignore these small companies."

As a result of the lack of clarity surrounding AI adoption pathways and risk ownership, SMEs are frequently forced into a false binary: either take on disproportionate risk or opt out of AI adoption altogether. The lack of onboarding infrastructure exacerbates this dilemma by compounding uncertainty and limiting the pathways through which SMEs can responsibly initiate their AI journey. One participant raised the issue of who ultimately bears the liability: when SMEs borrow or buy AI systems, who manages the risk if something goes wrong? These ambiguities create hesitation, widening the gap between intention and actual adoption.

With SMEs making up over 99% of enterprises in Southeast Asia, their exclusion from AI governance is not a theoretical concern; it is a systemic vulnerability. If AI governance is not designed to include them, it cannot be said to serve the public interest. As one participant put it starkly, "if SMEs fail, the economy fails". Support is thus not optional; ensuring that SMEs are structurally included in governance design is a foundational requirement for a functional AI ecosystem.

It is essential to recognise that the challenge is not that SMEs are inherently incapable of adopting AI responsibly. On the contrary, participants noted that SMEs are often more agile, more experimental, and more responsive to context than larger firms. The problem is that the supporting infrastructure, guidance, tools, incentives, and implementation support have not been designed or scaled to match the typical operating realities of SMEs. Without intentional support from more experienced players in the AI ecosystem, their agility is stifled by structural disadvantage.

Throughout the roundtable, participants suggested solutions, including the potential value of "jumpstart programs" and sector-specific support mechanisms, to help SMEs responsibly get started with AI. These were framed not as optional extras but as essential features of an inclusive governance environment. Participants also noted a growing recognition of the need to structurally include SMEs in the broader AI ecosystem. Suggestions included educational support and onboarding initiatives from financial institutions, peer-led learning networks within

professional sectors, and the systematic sharing of best practices across industries. However, more can be done, particularly in the domain of AI governance, where persistent ambiguity continues to hinder SMEs' ability to adopt AI responsibly.

Recommendation 2: Design the AI Ecosystem Around SME Enablement

For AI Ecosystem Anchors: Regulators, Banks, Tech Companies, and Sector or Industry Associations

AI Governance systems should be designed from the outset with SMEs in mind, not treated as an afterthought, retrofitted into ecosystems built around large firms, or left to navigate regulatory expectations on their own. This is the essence of the support-by-design principle: if Responsible AI is a shared societal goal, then AI governance must be structured to enable SME compliance without incurring prohibitive costs. As one participant put it, "We don't want to stop SMEs from developing. We want to give them a safe environment in which they can do this... create the governance [structures] that will help them to implement this safely... so that if anything goes wrong, they have done the best they can." AI startups must also be part of this conversation. These firms often operate under similar constraints to SMEs, particularly in terms of survival pressures and limited access to legal guidance. As one participant explained, "For startups and SMEs, survival is their day-to-day concern..."

Accountability for the responsible adoption of AI by SMEs must thus be shared across the ecosystem. Participants strongly highlighted that SMEs should not be expected to bootstrap AI Governance within systems built for organisations with far greater legal, technical, and financial capacity, especially since AI Governance is generally not the primary goal of businesses. One participant succinctly described the situation: "...to a business, these are externalities." A healthy AI governance ecosystem is thus one that distributes responsibility across institutional players. Regulators, banks, tech firms, and industry associations each play a role, not only in setting expectations but also in lowering barriers to adoption. This involves sharing technical knowledge, as well as co-developing specialised AI governance tools tailored to the business realities of the sector. As one participant from a bank asked, "Is there a white space for banks...government, and also players [at tech companies] to come together to support our SME community?"

This is not merely a call for ad hoc assistance, but for a more structural approach, one that integrates SME enablement into the foundations of the AI governance ecosystem. Again, as one participant put it starkly, "If SMEs fail, the economy fails."

AI Governance Gap: SMEs Are Expected to Adopt AI in Ecosystems Not Designed to Support Them

Recommendation: Regulators, Banks, Tech Companies, and Industry Associations—Design the AI Ecosystem Around SME Enablement

Support for SMEs must be by design and not incidental. A healthy AI governance ecosystem distributes responsibility across multiple actors:

1. **Regulators** play a key role by providing context-aware guidance tailored to different use cases.
2. **Banks** can support trusted onboarding processes and help drive learning across the sector.
3. **Tech companies** contribute by opening up access to tools, infrastructure, and governance scaffolding that smaller players can plug into.
4. **Industry associations** help localise standards, build sector-specific communities of practice, and organise stakeholders to co-develop specialised AI governance tools specific to the sector, thereby ensuring governance remains actionable.

Recommendation 3: Lead Real-World Pilots of AI Applications and Governance

For SMEs

SMEs are naturally positioned to iterate quickly. Their small size, mirroring Singapore's advantage, promotes agility and on-the-ground business context-awareness, making them ideal test beds for responsible AI adoption. Yet moving beyond experimentation remains a persistent challenge. While many SMEs are eager to pilot AI solutions, participants noted that scaling these efforts, whether across different business units or in sustaining them over time, can be difficult due to capacity constraints.

Importantly, while both large organisations and SMEs face obstacles when scaling AI, the challenges SMEs encounter are specific. Their challenges stem from resource constraints, particularly a lack of technical know-how in getting started, compounded by the difficulty of navigating an AI ecosystem often designed by and for larger firms. These challenges encompass cost, access to talent, and infrastructure, including compute resources. As one participant explained, "It's not that they are [overly] concerned with governance… They just don’t really know how to get started. And they don’t have the resources to get started."

This early-stage energy, however, should not be undervalued; it is a strategic asset that governance systems should actively support. As one participant noted, SMEs are often more responsive to local context and more willing to try new approaches than larger institutions. Yet

without the appropriate governance scaffolding, this potential remains underutilised. Indeed, participants frequently suggested building community-based mechanisms where SMEs can share both successful and unsuccessful case studies. As one participant put it, "We need that community of practice. We need the sort of horizontal sharing of good and bad experiences that will encourage experimentation, innovation, and learning from best and worst practices in this space." Another participant emphasised the importance of sector-based networks, noting, "Sector-based communities where you can share knowledge and you share your fate. Something goes wrong, we're all screwed. So having the knowledge, having the use cases, having the examples, case studies out there, so that once we do want to engage in it, it's easy to do so."

This suggestion is directly tied to the earlier recommendation of building a learning pipeline alongside a repository of successful AI governance use cases. A community repository greatly expands shared resources, ensuring economies of scale in pilot testing. Indeed, with SMEs playing such a significant role in the economy, success in AI adoption among them could generate broader network effects across the economy. With active iteration, perhaps SMEs can also account for a significant percentage of AI use cases. Through shared repositories alongside cross-sector learning, the next AI success story might very well come from an SME. With structured pilot testing and shared infrastructure, perhaps SMEs will not only represent 99% of firms in the broader Southeast Asian economy, but also 99% of their AI use cases and applications.

AI Governance Gap: SMEs Are Expected to Adopt AI in Ecosystems Not Designed to Support Them

Recommendation: SMEs—Lead Real-World Pilots of AI Applications and Governance

1. **The SME Opportunity: Position SMEs to lead AI use case experimentation**
   Leverage SME agility as a strategic asset. Their proximity to operational realities and capacity for fast iteration make them ideal early-stage innovators. AI governance systems should support SMEs not as an afterthought, but as front-line actors in experimentation.

2. **Led by Experience, Anchored by Expertise: Building a Community of Practice Around AI Adoption and Governance**
   Support cross-sector sharing by mobilising banks, regulators, and industry actors to foster horizontal learning. A vibrant community of practice, led by SMEs and anchored by experienced players, lowers the barrier to entry and aligns implementation with real-world experience.

3. **Develop a community repository of successful and unsuccessful AI use cases.**
   Build shared infrastructure for AI adoption and governance learning. Establish a shared knowledge base to accelerate adoption, minimise duplication, and document lessons

learned. A living repository allows SMEs to build *from* and *on* each other's experience, scaling good governance through network effects.

AI Governance Gap 3: Domain Experts (Non-tech Corporates) Are Excluded from AI Operations

AI pipelines today remain highly technical in nature, privileging model architecture, data processing, and optimisation at scale, often at the expense of contextual relevance. This design orientation sidelines domain experts who understand operational realities but lack access to meaningful intervention pathways. Yet in most organisational contexts, understanding how an AI output translates into business impact is inseparable from domain knowledge. Whether it's a risk assessment, medical recommendation, or policy flag, human judgment is required to guide or adapt outputs based on real-world conditions. However, because most AI pipelines are optimised for technical performance rather than usability, there are few accessible intervention points where non-technical experts can engage or assert oversight. The challenge is especially acute in settings where context matters at scale or in real time, such as finance, law, or healthcare. In the absence of intentional pathways for intervention, organisations lacking deep technical capacity are left to ask: where, when, and how can their human judgment be effectively integrated?

These risks are especially pronounced for SMEs, who often face an all-or-nothing AI proposition: either adopt off-the-shelf systems wholesale, with all the attendant risks, or forego AI adoption entirely. Unlike larger firms that can absorb integration challenges or tailor models in-house, SMEs are typically expected to adapt their business processes to fit off-the-shelf AI systems, not the other way around. This not only raises the cost of adoption but also introduces uncertainty into business outcomes. Without clear pathways to insert domain-specific judgment, SMEs lack the means to guide or constrain AI outputs in line with their operational realities. The result is a breakdown in trust, particularly with general-purpose technologies like GenAI, whose capabilities are broad but whose behaviours are often unpredictable and may not be aligned with sector-specific needs.

This structural omission has deep implications for AI governance. If AI governance is contextual, as most AI frameworks now acknowledge, it cannot be top-down alone. It must incorporate input from those who are closest to operational impact. The absence of clear pathways for domain expertise creates an epistemic gap: without insertion points for human judgment, there is no meaningful way to observe how AI systems behave in context beyond contrived benchmarks, nor to intervene when they behave undesirably. This lack of observability and responsiveness renders AI governance a theoretical exercise, grounded in principles but disconnected from business impact or outcomes. Ultimately, overly technical pipelines blunt the precision with which AI governance measures can be applied, leaving businesses with regulatory expectations they cannot operationalise along with risks they cannot meaningfully manage. All while relying on checkbox-style governance to keep mounting externalities at bay.

Recommendation 4: Insert Human Experts Here – Embed Organisational Domain Knowledge Into AI Workflows

For Organisations Adopting AI

AI governance should evolve from being viewed as “layers” added on top of AI development, such as in end-stage audits or reactive checklists. Similar to the first recommendation on translation workflows, this recommendation addresses the need to make governance tangible within day-to-day AI use. While translation workflows focus on interpreting external governance frameworks and aligning them with internal operations, this recommendation focuses on embedding organisational knowledge, including legal, operational, and product-specific knowledge, directly into AI workflows. If domain expertise is not embedded within AI workflows, it becomes challenging to identify context-specific risks or effectively adapt outputs to operational realities. Embedding human judgment throughout the AI lifecycle, including during design, deployment, and refinement, is crucial to ensuring that systems accurately reflect operational realities, not just engineering goals. Non-tech corporates or business domain experts contribute more than organisational knowledge; they provide the interpretive grounding needed to align model outputs with organisational intent. This includes governance perspectives, especially since risk mitigation relies on understanding the effects or outcomes of AI use; thus, this issue should not be entrusted solely to engineering teams. Achieving this shift requires rethinking how AI teams are composed and how workflows are structured, so that organisational, humanistic, and technical knowledge are integrated from the outset, rather than retrofitted into workflows that silo perspectives as a structural blind spot.

This is where AI governance expertise, which lies at the intersection of AI systems and organisational risk, can play a defining role. Rather than treating governance as a checklist at the end of a project, these experts can adapt AI workflows to align with governance objectives and organisational priorities. Their role is not only to interpret risk; it also involves surfacing where and how AI limitations should shape deployment decisions. A critical function of this work is designing intentional intervention points, moments where human judgment can observe, guide, adjust, or constrain AI outputs within the appropriate timeframe relevant to the business context. These may occur during model training, validation, post-deployment testing, or in real-time. Without such mechanisms, governance remains abstract; the opportunity to make AI both accountable and adaptive is thus lost. Such intentional intervention points can take the form of AI governance tools, where tools act as the interface through which domain experts can flag concerns, adjust outputs, or pause automation. By providing controlled points of intervention, tools transform judgment from ad hoc review into a formalised part of the AI lifecycle.

As one participant noted, "one of the issues raised... is that definitions for AI are still unsatisfactory. Not everyone can understand why the ambiguity is a problem..." This lack of clarity reinforces the importance of people who can serve as translators between technical systems and regulatory or business objectives, as well as embed domain expertise where

necessary. Another participant from law enforcement shared a scenario in which AI was used for analysis and later challenged in court. Because the analyst deferred entirely to the AI's recommendation without understanding how it arrived at its conclusion, the court rejected the evidence as unreliable. The participant remarked, "It's going to be us who have to go to court and explain why AI did what it did..." This exemplifies the need for structured mechanisms where human judgment is not only present but meaningfully embedded throughout the AI lifecycle, especially when humans ultimately bear accountability for the outputs of AI. Such obstacles in interpretability stem from the well-known black box nature of AI. While the black box nature of AI remains a fundamental challenge that may ultimately require advances in core AI research, embedding domain expertise within AI development processes enables human judgment to come into play in practical, immediate ways. This includes leveraging existing technical levers, such as fine-tuning and reinforcement learning with human feedback, as well as other technical methods that can influence AI outputs. This approach enables not only the assessment of the overall feasibility of AI deployment in particular use cases but also the scrutiny of specific outputs and the handling of issues as they arise. For SMEs, which may lack in-house capacity to deploy these methods directly, ecosystem anchors can play a key role in lowering the barrier to entry. Community of practice networks can also help SMEs identify appropriate tools, access shared resources, and adapt successful oversight strategies developed by peers. This shared infrastructure enables SMEs to apply governance mechanisms that would otherwise be out of reach, making real-time human judgment both accessible and actionable.

AI Governance Gap: Domain Experts Are Excluded from AI Operations

Recommendation: Organisations Adopting AI—Embed Organisational Domain Knowledge Into AI Workflows

1. **Embed Domain Knowledge, Including AI Governance, into AI Workflows by Creating Specific Intervention Points for Human Judgment**
   Integrate domain and governance experts across the entire AI pipeline. Structure deliberate moments within existing BAU workflows, such as during training, testing, or real-time use, where human judgment can observe, guide, or, if necessary, constrain system behaviour.
2. **Use Existing Technical Levers to Make Human Judgment Actionable**
   Domain experts can leverage tools like fine-tuning, reinforcement learning with human feedback, and prompt-based controls to operationalise human oversight and influence AI outputs, even without full model interpretability. SMEs can draw on ecosystem anchors and communities of practice to access tools, resources, and oversight strategies that make human judgment in AI workflows practical and actionable.
3. **Human Intervention Can Constrain the Implications of the Black Box Nature of AI**
   Even when AI systems are fundamentally opaque, embedding human judgment

throughout the AI workflow, for example, in the form of AI governance tools, which form the interface between domain experts and AI workflows, enables organisations to assess the feasibility of AI deployment, scrutinise outputs, and make risk-informed decisions grounded in context.

AI Governance Gap 4: AI Governance Is Often Framed as a Constraint

AI governance is still largely framed as an obligation or constraint, a risk-mitigation tool that organisations must comply with, rather than as a catalyst that can be actively leveraged to shape AI innovation. This framing carries significant consequences. How corporate stakeholders discuss governance influences how AI is adopted, designed, and who feels empowered to participate in its development. When governance is communicated as a compliance-driven obligation, it discourages engagement, especially from non-technical teams. As one participant put it, "How can governance support AI adoption? That's not something you normally associate together… but there's a lot of meaning in that." Another described how ambiguity around what counts as "AI" leads to hesitation, disengagement: "When they don't know the implications of actually calling something AI, then you become very scared." Framing matters. It shapes perceptions of trust, accessibility, and ownership, ultimately determining whether governance is understood as a *catalyst* or a *constraint*.

This framing doesn't just influence understanding; it shapes participation. It determines which stakeholder engages, how they contribute, and what actions they feel authorised to take. A compliance-based narrative tends to promote caution, resistance, encouraging teams to "slow things down" rather than build governance into workflows. A catalyst-based framing, by contrast, positions governance as part of the AI adoption process, creating space for proactive, context-aware measures. As one participant noted how that can be done, "...governance has a better shot at supporting adoption and innovation when it's structured according to use cases." This perspective highlights how structural changes to the way AI governance is framed, especially when governance is framed in contextually meaningful and operational terms, can influence AI adoption outcomes. Another emphasised the chilling effect of ambiguity, explaining, "It kind of inhibits innovation, courage… when they don't know the implications of actually calling something AI, then you become very scared." One participant further said, "This is not a [just] compliance activity…this is a journey we are all on together." When AI governance is communicated as a shared organisational function rather than institutional risk control, it invites broader participation.

This challenge is further complicated by the fragmentation of AI governance across jurisdictions. Divergent regulatory narratives, whether rooted in the EU's precautionary stance, the US's market-first ethos, or China's state-centric approach, amplify the perception that governance is inconsistent and externally imposed. One participant described it as different versions of an all-encompassing "blanket approach," highlighting the difficulties arising from the interoperability of different jurisdictions. For non-tech corporates operating across markets, this ambiguity erodes confidence and reinforces the sense that governance is something to be avoided, rather than engaged with.

Clear internal framing is therefore crucial. AI governance experts involved in the translation workflow can help construct a coherent governance narrative, reconciling external regulatory expectations while ensuring internal clarity. This enables organisations to navigate governance proactively rather than reactively.

Recommendation 5: The Role of AI Governance as a Catalyst for Collective AI Innovation

For Governments, Regulators and AI Governance Experts

AI governance can serve as the connection that enables collective AI innovation across the organisation. When viewed as a catalyst, the governance function can play the role of establishing the internal framework, alongside the structures and roles, necessary for responsible and coordinated progress in AI:

- Business domain experts are assigned specific roles within AI workflows.
- Governance-related priorities such as safety, security, ethics, and alignment have a clear operational base.
- Technical teams can build with confidence, assured that their work aligns with organisational goals, regulatory requirements, and societal expectations.

This "whole-of-organisation" approach enables AI innovation to be forward-looking in terms of technical capabilities while remaining grounded in the operational context. It is mindful of risk and potential harm, impacting both the organisation and society at large. Without this reframing, AI innovation tends to be fragmented: technically advanced yet misaligned with the realities and responsibilities of the organisations deploying it.

AI Governance Gap: AI Governance Is Often Framed as a Constraint

Recommendation: Governments, Regulators and AI Governance Experts—Frame AI Governance Not as a Constraint, but as a Catalyst of Collective AI Innovation

1. **Framing Matters: Position AI Governance Not as Constraint**
   Language shapes how governance is perceived and who engages with it. Moving away from compliance-first narratives enables governance to be understood as a practical support structure for adoption, rather than a barrier.
2. **Use AI Governance as the Bridge for Inclusive Whole-of-Organisation Participation**
   When framed correctly, governance becomes the mechanism that connects technical teams, domain experts, and leadership, ensuring that everyone has a meaningful role in shaping how AI is deployed.

3. **AI Governance as the Catalyst of Collective AI Innovation**
   Governance does not hinder innovation; rather, it grounds it, contextualises it, promotes shared responsibility, and ensures alignment. It ensures that AI innovation progresses with organisational intent, context sensitivity, and adherence to societal and regulatory standards.

Conclusion: A Vision for AI Governance

Who Needs to Do What: A Shared Mandate for Collective AI Innovation	
Stakeholder	**Role in Enabling Collective AI Innovation**
Organisations Adopting AI	Establish translation workflows to align governance with day-to-day operations. Embed legal, operational, and product-specific knowledge into AI pipelines. Build internal alignment and contribute to shared governance learning through real-world application.
Corporates (Domain Experts)	Provide grounded, context-specific judgment in the design and use of AI systems. Ensure that governance reflects sector realities, customer expectations, and business operations. Act as accountability anchors where AI impacts real-world decisions.
SMEs	Act as first-movers in AI experimentation. Lead development of context-specific use cases. Share learning through open repositories. Drive grassroots innovation grounded in real operational constraints.
Ecosystem Anchors (e.g. regulators, tech firms, industry associations)	Support SMEs by anchoring communities of practice. Provide access to governance scaffolding, infrastructure, and domain expertise. Help translate frameworks into usable tools and localise adoption strategies. Contribute to the co-development of contextualised AI use cases and shared repositories.
Governments and Regulators	Provide context-aware, use-case-driven guidance. Coordinate cross-agency efforts. Shape the public narrative by framing AI governance as an innovation catalyst. Promote long-term coherence across jurisdictions and sectors.
AI Governance Professionals	Serve as integrators across tech, product, legal, and compliance teams. Design translation workflows that bridge high-level governance with operational realities. Navigate regulatory complexity and embed governance into daily operations.

The AISG roundtable held in collaboration with Deloitte Singapore on *How can governance support AI adoption among non-tech corporates?* and *How can industry and corporations navigate the proliferation of AI governance across multiple countries and regions?* surfaced a clear consensus: AI governance cannot remain abstract, top-down, or external to operations. For governance to enable adoption and build trust, it must be embedded in practice, contextual to each sector, and co-owned across institutions. Across the discussion, participants highlighted four key governance gaps: First, that existing frameworks often fail to translate into operational guidance, leaving organisations with expectations they cannot act on. Second, that SMEs are expected to adopt AI in ecosystems not structurally designed to support them. Third, that AI pipelines exclude domain experts, making systems technically sound but contextually brittle.

And fourth, that governance is often framed as a constraint rather than a catalyst, discouraging innovation and sidelining broader participation. Taken together, these gaps point to a critical insight: the prevailing paradigm, where AI governance is perceived as just compliance or a brake on innovation, must be reimagined.

The discussion did not stop at analysing gaps; participants actively proposed solutions. A core theme throughout was the need to move away from governance that is externally imposed, toward governance that is developed with and through organisations. Participants emphasised that effective AI governance should be integrated into daily operations, rather than added as an extra layer afterwards. This shift requires designing governance that is operationally relevant, context-aware, and inclusive across organisational roles. To that end, participants surfaced five key recommendations: first, to develop a translation workflow that bridges high-level governance principles and real-world AI workflows; second, to embed domain experts into AI pipelines by creating deliberate intervention points for human judgment; third, to empower SMEs through governance ecosystems intentionally structured to support them; fourth, to embed organisational domain knowledge into AI workflows; and finally, to reframe governance as the catalyst that enables collective, responsible AI innovation. As the table on "Who Needs to Do What: A Shared Mandate for Collective AI Innovation" shows, AI governance leading to collective AI innovation is a shared endeavour.

Ultimately, the answer to the question of how governance can support AI adoption among non-tech corporates begins with their inclusion. AI governance cannot be meaningfully developed or deployed without the active participation of non-tech corporates, including those from both large organisations and SMEs. These stakeholders are not only domain experts but are often accountable for the real-world outputs of AI. Yet, today, AI governance is often still treated as a technical, legal, or compliance issue. As this roundtable has made clear, it is equally a question of participation, partnership, communication, culture, and institutional design. While technical stakeholders will continue to advance the frontier of AI research, non-tech corporates have a unique opportunity to lead where AI truly matters: in its application, its impact on people's lives, and its alignment with the values we live by. Governance should thus not only be seen as a cost of AI adoption, but rather, like participants envisioned, as a catalyst for collective AI innovation, making it usable, inclusive, and safe. If we collectively view governance as a bridge rather than an impediment, we will design AI systems that incorporate human expertise at the right juncture, thereby entrenching our shared humanistic perspectives in a technology often portrayed as being misaligned or at odds with us. The question then is no longer whether governance will support AI adoption; it is how deliberately and how collectively we are willing to shape our shared future with the empowerment of AI.

Acknowledgements

This roundtable was founded via a charitable grant from Google.org, as part of Google's Digital Futures Project, and hosted by Deloitte Singapore.