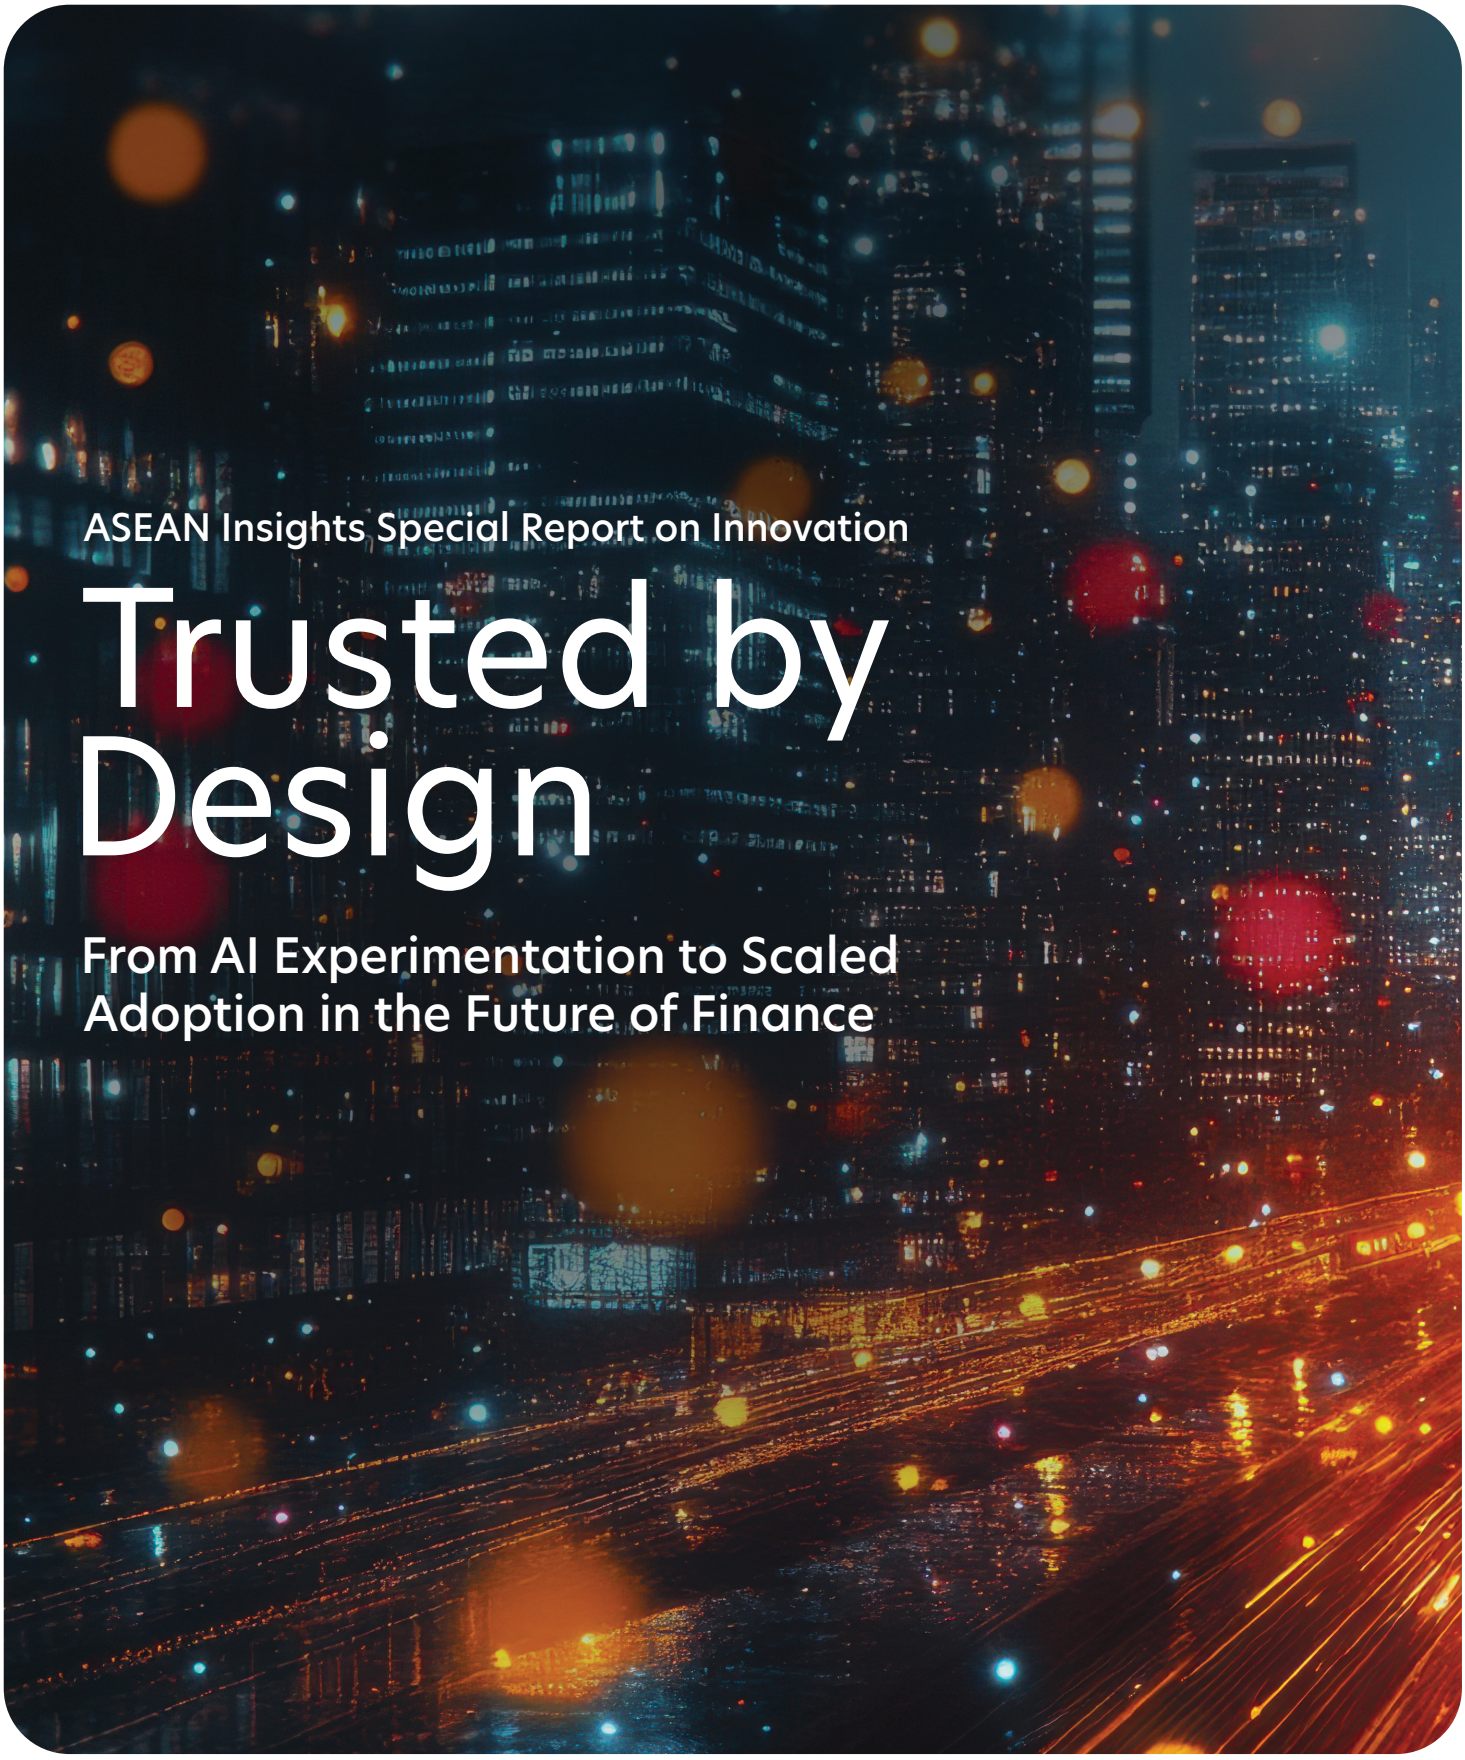
UOB



ASEAN Insights Special Report on Innovation

Trusted by Design

From AI Experimentation to Scaled Adoption in the Future of Finance

November 2026

Right By You

Table of Contents

Foreword by Lee Zhu Kuang, Head of Innovation Group, UOB Pg 3

Section 1: Two shifts define the moment Pg 4

Section 2: AI as a core utility for finance Pg 5

Section 3: Three tensions challenging trusted AI adoption Pg 6

Section 4: What future-ready organisations are doing Pg 8

Section 5: The UOB perspective – trusted adoption in practice Pg 9

Section 6: Conclusion Pg 12



Foreword

Innovation is not defined by technology alone. It is measured by an organisation's ability to adapt to change, create value, and deliver better outcomes for customers, businesses, and society.

Today, we are entering a new phase of transformation. Advances in digital technologies are reshaping industries, accelerating decision-making, and creating new opportunities for growth. Yet the differentiator is no longer access to technology. It is an organisation's ability to harness these capabilities responsibly, scale them effectively, and translate them into meaningful business outcomes.

At UOB, our experience has reinforced a simple belief: transformation begins with strong foundations. Trusted data, sound governance, skilled talent, and a culture of continuous learning are essential to turning innovation into sustained impact. Equally important is collaboration across the broader ecosystem, bringing together industry, regulators, technology partners, and academia to address shared challenges and unlock new opportunities. Not just for ourselves but also for ASEAN and for good.

As organisations seek to move from experimentation to enterprise-wide adoption of emerging technologies, trust becomes increasingly important. Trust is not a constraint on progress. It is what enables innovation to scale with confidence, resilience, and accountability. Those that successfully integrate innovation with trust will be best positioned to navigate uncertainty, strengthen competitiveness and create long-term value.

This report brings together perspectives from leaders across industry, government, and academia on how organisations can build for the future while remaining grounded in the principles that matter most. While technologies will continue to evolve, one conviction remains constant: sustainable progress is achieved when innovation is trusted by design.

Thank you to the following individuals and teams for contributing perspectives and insights that informed this report:

- Chia Hock Lai (Responsible Fintech Institute)
- Professor Louis Phee (Nanyang Technological University)
- Mike Sim (Independent Contributor)
- Chang Sau Sheong (Government Technology Agency of Singapore)
- Sutowo Wong (Temus)
- Varun Mahindru (Visa)
- UOB FinLab & Ecosystems Partnerships
- UOB Innovation Group



By Lee Zhu Kuang,
Head of Innovation Group, UOB

From experimentation to adoption. The proof-of-concept economy has produced a proliferation of pilots. Most work in isolation. But only a few have been scaled into how the business actually runs. The bottleneck between a promising pilot and organisational impact is rarely the model. It is the process redesign, workflow integration, people readiness and organisational alignment required to move from what works in a small experiment to what changes how work happens.

From innovation projects to core capability

For a long time, AI in financial services was framed as an innovation initiative — a discrete project with its own justification, ROI case and separate leadership. That framing is increasingly out of step with how AI is being consumed. Basic AI is now embedded in almost every enterprise software subscription and every productivity tool. It is becoming a mainstream business capability rather than an experimental innovation. Institutions that still treat all AI as innovation will find themselves paying twice — once for capabilities that are now baseline, and a second time for the deployments that genuinely justify a separate business case.

A related architectural point

Cloud computing began as an innovation initiative and later became a procurement exercise. The internet began as a strategic differentiator and over time became a standard utility. Enterprise software followed the same arc. AI is following a familiar path from innovation to utility, but at a significantly faster pace. Institutions that see this pattern and adapt their operating discipline accordingly will have a lasting advantage over those still debating whether the AI wave has arrived.



Why baseline AI is becoming infrastructure, and what that changes

The most consequential shift in AI in the last two years is a quiet one. Basic AI capability has moved from a differentiated investment into a baseline utility.

Many major productivity platform now include some form of embedded AI. Every enterprise software renewal is bundled with an AI assistant. Generative AI tools and capabilities for coding, drafting, summarising, transcribing and analysing sit inside tools that were bought for other reasons. In a growing number of finance functions, AI assisted knowledge retrieval, drafting and decision support are becoming as ubiquitous as spreadsheets. According to the UOB Business Outlook Study 2026, 65% of businesses surveyed have deployed AI, yet only 15% have reached advanced AI capabilities, underscoring the gap between adoption and maturity.

The pattern is familiar to anyone who watched cloud, the internet or enterprise software move from strategic differentiator into procured utility. AI is following the same arc, faster.

This shift changes how organisations should think about AI investment. The question is no longer simply whether to invest in AI, but what type of AI capability is being considered.

Two categories of AI investment are emerging

The current landscape indicates two categories that have very different economics, different governance requirements, and very different questions to answer.

Foundational AI capability:

This is the AI equivalent of cloud storage or internet connectivity. Copilot licences, productivity assistants, embedded generative AI in enterprise software, and general purpose knowledge tools. The value case here is not a discrete ROI. Having these capabilities is now baseline and not having them puts an organisation at a productivity, retention and talent disadvantage. Institutions that continue to treat this layer as a series of separate business cases, each requiring individual justification, underestimate the opportunities they can provide.

Strategic AI investment:

Agentic workflows that handle multi-step business processes autonomously. AI platforms embedded into core customer servicing, credit decisioning, fraud prevention or relationship management. Bespoke models built on proprietary data. These deployments genuinely require business case discipline. They involve real capital, significant integration risk, detailed governance requirements and real change management. Applying the same investment approach as foundational productivity tools may therefore understate their complexity and risk.

The strategic risk is conflating these two categories. Applying the same investment discipline to foundational AI as to strategic deployments can slow adoption and productivity gains. Conversely, applying insufficient scrutiny to strategic AI investments can increase enterprise risk. This calls for a clearer distinction between the two treating foundational AI as infrastructure that is procured and rolled out, and reserving intensive governance, business case scrutiny and change management for the strategic deployments where those disciplines create value.

> **If agents are the ones making purchases on behalf of humans, how does advertising work? Are you advertising to the human or the agent? The entire commercial model of digital finance is being renegotiated, and most institutions have not started the conversation.**
>
> **– Industry perspective**

Technology is not the primary bottleneck. Adoption is. Three tensions define what is actually hard about scaling AI responsibly in financial services. How organisations navigate them will determine which institutions define the next chapter of finance.

Tension 1 Speed vs Trust

The pressure to move quickly on AI is real. Every board wants deployment. Every competitive analysis argues that slower means falling behind. But the same pace that produces impressive demonstrations produces the failures that erode trust. When Apple Card faced accusations of gender biased credit limits, subsequent investigation found no discrimination in the algorithm. What harmed the brand was the absence of explainability. This points to a broader principle: explainability should not be viewed simply as a regulatory requirement, but as an asset for building trust.

Trust is doing serious work in AI adoption. Customers, employees and regulators are converging on a common question: "can I trust how your bank uses AI?" Organisations that treat trust as a compliance layer to be added later find that adoption stalls, because internal stakeholders, business units and customers cannot rely on the outputs. Organisations that treat trust as a design principle building explainability, oversight and accountability into deployment from the outset find that adoption accelerates, because confidence is embedded rather than argued for. This reframing matters. Governance and responsible AI are often positioned as constraints on innovation necessary evils that slow the pace of deployment. That framing is exactly backwards. When embedded by design, governance and responsible AI accelerate adoption because they create the trust, accountability and confidence necessary to scale AI safely and sustainably. The practical implication is significant: making governance visible showing customers, employees and regulators how AI decisions are made, what data was used, what oversight exists and what recourse is available becomes a source of competitive advantage rather than a compliance cost.

Trust in AI also needs to be considered at three levels simultaneously: at the individual level (the engineer, the advisor, the officer using the tool); at the organisational level (governance principles, ethics review, oversight cadence); and at the sectoral level (regulatory guardrails coordinated across markets). Most institutions currently operate meaningfully at only one of those three. That is not enough. Trust, in this framing, is not a brake on speed. It is the enabler of sustainable momentum. The practical principle that follows — and one that has guided UOB's own scaling work is straightforward: start with trust before scale. It is far cheaper and faster to build trust into an AI capability at design than to retrofit it after deployment has begun to fail.

> **Accountability for AI has to be designed at three levels. The individual using the tool, the organisation setting the policies, and the sector coordinating the guardrails. Most institutions currently operate meaningfully at only one of those three. That is not enough.**
>
> **— Industry perspective**

Tension 2 Experimentation versus Scaled Adoption

This is the tension the market has not yet fully reckoned with. Many organisations have run hundreds of AI pilots. But few have scaled a meaningful number of them into production and workflows for everyday use. The gap between "we ran the pilot" and "the business runs differently because of it" is where AI's promise most often disappears. The reasons are not primarily technical. The pilots work. What breaks down is everything around the pilot. The process is not redesigned to make room for the new capability; the workflow still routes work through the old paths. The people using the tool were not brought along; they revert to what they know.

> **The real cost in an organisation is not in deploying the AI agents. It is in the whole reimagination of the workflow around them. That is what most proofs-of-concept leave undone, and that is why they do not scale.**
>
> **— Industry perspective**

The organisational alignment - leadership sponsorship, incentive structures, cross-functional coordination is not there. Governance was designed for the pilot, not for the scaled deployment. So the deployment stalls, and the pilot is quietly deprioritised. Scaling AI responsibly requires organisations to treat adoption as a distinct discipline with its own methodology, leadership and metrics. AI industrialisation, in this sense, is fundamentally about operating model — the governance structures, platforms, talent strategies and workflow redesign required to move AI from isolated use cases into enterprise-wide capability. It requires process redesign, not just process automation. It requires investment in the everyday user experience of AI, not just the capability of the underlying model. And it requires governance that enables adoption rather than constrains it — embedded upstream in the design of the deployment rather than layered on afterwards as approval gates.

Tension 3 Technology Capability versus Human Readiness

The productivity case for AI is empirically clear. The human case what AI means for the people who work in and are served by financial institutions is unresolved and largely absent from the public conversation.

Boston Consulting Group's latest workforce research reinforces that people are fundamental to AI adoption, with around 70% of AI value linked to 'rethinking the people component', including upskilling and structured learning. According to the UOB Business Outlook Study 2026, ASEAN SMEs adopting AI cite workforce anxiety as one of the top three barriers to deeper adoption. The barrier is not that the technology is missing. It is that the people it is being deployed onto are not ready not because they are unwilling, but because their skills, workflows and confidence have not been built in step.

As one innovation leader observed, "If AI can free the doctor from typing notes to actually be present with the patient, that itself is profoundly impactful. We need to see innovation as a source of pride and trust, not as a source of job replacement."

Underneath the workforce anxiety numbers sit specific human questions that most institutional AI strategies do not explicitly answer. Employees are often less concerned about the technology itself and more concerned with: How will my role evolve? Will my skills remain relevant? How can I use AI responsibly? Addressing these questions requires more than a communications programme. It requires trust, structured capability building, visible leadership commitment, and a culture that encourages experimentation and learning.

Two principles from UOB's own experience translate that requirement into a practical stance. First, start with trust before scale make governance, accountability and responsible AI practices visible so that confidence is built before the wider rollout, not after. Second, equip every employee, not just technical teams enterprise-wide AI literacy enables employees across every function to identify opportunities and create value. One of the most significant institutional inflection points in an AI journey occurs when employees move from simply consuming AI tools to building solutions themselves from users of AI to builders with AI.

The organisations moving fastest and most sustainably on AI are the ones investing in workforce readiness, leadership readiness, skills development and change management as first order activities — not as HR appendices to the technology programme. Augmentation, not replacement, is not a slogan. It is the design principle that determines whether adoption is sustainable.



From theory to observable practice

Across the more thoughtful adopters in financial services, four patterns are emerging. They are less a checklist than a set of observable practices — the things that seem to be common to the organisations moving from AI experimentation to embedded, scaled adoption.

Pattern 1 Embedding AI into workflows, not adjacent to them

The first shift is architectural. AI is being placed inside the workflow the business already runs, rather than sitting in a separate tool that employees have to remember to use. The distinction sounds trivial but has significant consequences. When AI is adjacent — a separate tab, a separate app, a separate approval flow — adoption depends on habit change, and habit change is where most rollouts fail. When AI is embedded in the workflow, the incremental cognitive cost of using it approaches zero. That is what turns a promising pilot into a scaled capability. Practically, this means treating each AI deployment as a workflow redesign challenge first and a technology implementation second. Organisations should design AI around how work is actually done, integrating it into the tools and processes people already use. This reduces friction, encourages adoption and helps turn promising pilots into capabilities that can scale.

Pattern 2 Building trust into deployment from the design phase

The organisations getting adoption right are treating trust as an upstream design requirement rather than a downstream governance overlay. Explainability, oversight mechanisms, error handling, accountability lines and user recourse are specified alongside the model, not added afterwards to satisfy compliance. Making these design decisions visible internally to the employees who will use the AI and to the customers it will serve is what turns governance into an adoption accelerator rather than a friction point. The operating principle that follows is what UOB has learned to call "start with trust before scale": embed responsible AI from the outset, make it visible, and let confidence do the work of adoption. That is what shifts governance from a bottleneck into an enabler.

> **In surgical robotics, when something went wrong the doctor was accountable — simple. With AI in the picture, that line of accountability changes. The question of who owns the decision becomes the design question itself.**
>
> **— Industry perspective**

Pattern 3 Investing in people alongside technology

Serious AI adopters are investing in the human side of the deployment at least as heavily as the technical side. Training goes beyond tool familiarity into genuine capability building. Career paths are redesigned so that AI augmented roles have credible progression. Leadership development ensures managers know how to run teams whose work is being transformed. Change communication treats employees as participants rather than recipients of a new tool.

Critically, the organisations that scale AI most sustainably extend enterprise wide AI literacy — they equip every employee, not just data scientists and engineers, with the confidence and vocabulary to spot opportunities and put AI to work in their own functions. The most significant institutional inflection point in an AI journey happens when employees move from consuming AI tools to building solutions themselves. That transition — from users to builders — is when adoption compounds. It cannot be delegated to the technology function; it has to be led from the top and supported enterprise wide.

Pattern 4 Working through ecosystems, not in isolation

The organisations most effective at scaling AI adoption are those that recognise no single institution can build all the required capability alone. They are working through ecosystems — with regulators, technology partners, academic institutions, industry bodies and SME communities. Regulators become co-designers of adoption pathways rather than external gatekeepers. Technology partners become long term capability builders rather than one off vendors. Academic institutions provide independent voice and validation. Industry bodies convene shared learning. SME networks become the testing ground and feedback loop for what works in practice.

This ecosystem posture is one of the more reliable predictors of whether an AI adoption strategy will scale. Institutions that treat AI adoption as a solo enterprise consistently underdeliver against institutions that treat it as a networked endeavour. The next phase of financial services will not be dominated by the institution with the most impressive internal AI stack. It will be shaped by the institutions that have built the ecosystem relationships to scale adoption safely across their markets.

Scaling Trusted AI with IMPACT

UOB's AI vision is to scale Advanced AI responsibly to deliver enterprise value and enable an AI-ready workforce. IMPACT translates this vision into six strategic themes that connect innovation & research, enterprise value, modern platforms, accountable deployment, culture and talent.

I

Innovation & Research

Build AI leadership and strategic differentiators through AI innovation, research and strategic ecosystem partnerships.

M

Measurable Value

Deliver tangible economic value to the business through revenue growth, improved customer engagement, strengthened risk management and productivity gains.

P

Platforms Modernisation

Modernising enterprise AI, data and knowledge platforms to scale trusted intelligence.

A

Accountable AI

Embed responsible AI and governance into every stage of the AI lifecycle to ensure safe and ethical AI deployment.

C

Culture

Ignite a culture of experimentation, curiosity and continuous improvement across the Bank.

T

Talent

Empower an AI-enabled workforce through skills acceleration, talent acquisition and mobility, and assurance.

From framework to practice

The practical lessons that follow show how UOB puts IMPACT into practice: strong data foundations and modern platforms; responsible governance embedded from the outset; disciplined value measurement; enterprise-wide capability building; and ecosystem-led adoption across ASEAN SMEs.

Proof, not theory

The most useful thing UOB can add to the conversation on scaled AI adoption is not another prescription. It is the record of an institution that has been living the transformation from the inside for more than a decade, at enterprise scale, and simultaneously extending the same discipline to ASEAN SMEs through UOB FinLab. Neither strand is a case study to be pointed to when convenient. Both are working methods, and the patterns they surface are what the rest of the market is now trying to articulate.

The paper draws four sets of practical lessons from that experience: building the data foundations before the AI, scaling AI as an enterprise capability, measuring value rather than activity, and extending the same discipline to ASEAN SMEs.

A. Building the data foundations before the AI

At UOB, the AI journey did not start with generative AI. It started with data. More than a decade of deliberate investment in enterprise data architecture, governance, stewardship, quality management and analytics platforms produced what generative AI would later need: trusted, accessible, governed data available consistently across the organisation.

The specific investments included the Enterprise Data Architecture Governance (EDAG) framework, integrated data lake and warehouse capabilities, self-service analytics, and data quality frameworks. None of these were designed with generative AI in mind. All of them turned out to be the enabling infrastructure once generative AI could productively use them. The lesson is straightforward and unavoidable: before scaling AI, establish strong data foundations. Institutions treating data foundations as a downstream problem — to be tackled after the AI initiative has been announced consistently underestimate how long they take to build and how much they compound over time. The organisations moving fastest on AI today are the ones that were quietly investing in their data layer years ago.

B. Scaling AI as an enterprise capability

Building on those foundations, UOB has scaled AI through a multilayered approach. Today, approximately 30,000 employees across the organisation have access to Copilot capabilities. More than 30 Build Your Own Bot initiatives have been implemented. Over 300 AI and Data Analytics use cases have been deployed across customer engagement, risk management, fraud detection, AML/KYC, credit analytics, productivity and decision intelligence.

These numbers are evidence of what strong foundations, governance and workforce enablement translate into once you are actually scaling. Two operating principles have guided how that scaling was done and they are the ones worth taking away.

Start with trust before scale:

Responsible AI is embedded into innovation from the outset — not layered on afterwards. By making governance, accountability and responsible AI practices visible, the organisation builds confidence and encourages wider adoption. This is the practical form of the argument in Section 2: trust is the enabler of scale, not a friction cost against it.

Equip every employee, not just technical teams:

AI cannot be limited to data scientists and engineers. Enterprise-wide AI literacy and capability building enable employees across every function to identify opportunities and create value with AI. The most significant institutional inflection point occurs when employees move from simply consuming AI tools to building solutions themselves from users of AI to builders with AI.

C. Measuring value, not activity

The measurement discipline discussed in Section 2 is not an abstract framework. It is a live piece of institutional work at UOB. The bank is developing a bank-wide enterprise value measurement framework designed to consistently, transparently and auditably measure the economic value generated by AI and Data Analytics investments. The purpose is straightforward: to move AI value measurement from subjective estimation to evidence based, repeatable assessment across the three dimensions of revenue, cost and capacity, and risk.

The point of doing this work at institutional scale rather than at the level of individual projects is that it changes the conversation. Once AI value is measured consistently and audibly, boards, regulators and management teams can compare investments on a common basis. Business units stop making incompatible claims for their own initiatives. AI investment becomes a strategic capital allocation discussion, not a series of narrative pitches. This is the discipline that separates institutions treating AI as a permanent capability from institutions treating it as a rolling series of experiments.

D. Extending the same discipline to ASEAN SMEs

The second strand of UOB's experience is the work UOB FinLab has done with ASEAN SMEs, most publicly through the AI Ready Programme and the UOB SME Elevate Programme. Across cohorts, four lessons keep surfacing. They are the same lessons the enterprise transformation surfaces internally which is precisely why they matter.

Start with the business problem, not the technology capability. Every successful UOB FinLab AI adoption journey begins with a structured discovery process that identifies where AI can create genuine business value in a specific company's real workflow. Technology decisions follow the business problem.

Design for augmentation, not replacement:

AI is introduced to reduce manual and repetitive work, free time for higher value activity, and grow capacity not as a cost reduction lever. Businesses that frame AI as augmentation retain their people and grow. Businesses that frame it as headcount reduction see trust collapse and adoption reverse.

Build user confidence alongside technical capability:

Cohort learning, industry mentorship, hands-on coaching and staged proof-of-concept are used to build user readiness alongside technical capability. This is the SME equivalent of "equip every employee".

Combine technology with human support and ecosystem partnership:

AI adoption is not a technology project. It is a business transformation journey. Through UOB FinLab, businesses gain access not only to technology solutions, but also trusted advisory, peer learning and a wider ecosystem of partners that help them move from exploration to implementation. This combination of technology, human guidance and ecosystem support is what turns experimentation into real business outcomes.

Taken together, the enterprise transformation and the UOB FinLab work describe the same posture at two different scales. Data foundations first. Governance built in from the start. Every employee equipped, not just the specialists. Value measured against outcomes, not activity. Ecosystems used deliberately rather than left implicit. These are not innovations UOB is announcing. They are patterns UOB has been operating on for years — and they are, this paper argues, the patterns that will define trusted AI adoption across ASEAN in the decade ahead.

Trusted AI adoption as a core capability for the future of finance

This paper has argued that the next phase of value creation in financial services will not come from more AI experimentation. It will come from scaled, trusted, embedded AI adoption the kind that changes how organisations operate rather than what they demonstrate. It is also, honestly, the harder work. Experimentation is intellectually rewarding and organisationally cheap. Scaled adoption requires operating model change, data foundations, governance embedded from the start, workforce investment, value measurement discipline, and ecosystem partnership. It is the work many institutions are least prepared for.

The institutions that will define the next decade of finance in ASEAN are the ones investing in that work now. They are drawing a clear line between baseline AI utility, strategic AI investment, and the data layer beneath both. They are embedding AI into workflows rather than adjacent to them. They are treating trust as a design principle, not a compliance obligation. They are equipping every employee, not just the technical function. They are building the value measurement discipline that turns AI investment into a serious capital allocation conversation. They are working through ecosystems rather than in isolation. Above all, they are moving from experimentation to adoption with the deliberateness the moment requires.

UOB's own transformation, the decade of data foundations, the enterprise-wide adoption scaled across approximately 30,000 employees, more than 300 deployed use cases, the value measurement framework now being built, and the UOB FinLab discipline extended to ASEAN SMEs, is offered here as one institutional example among many that are possible.

The point is not that UOB has completed the journey. It has not. The point is that the shape of the journey is now visible, and the institutions that see it clearly and commit to it will be the ones that shape the future of finance in ASEAN and for good.

Trusted AI adoption is becoming a core capability for the future of finance. Institutions that scale adoption while retaining trust will define that future. That is what trusted by design means. That is the standard the moment demands - the patterns that will define trusted AI adoption across ASEAN in the decade ahead.



This report draws on UOB's experience and research, together with perspectives shared by leaders and practitioners across industry, government and academia during a UOB-convened expert roundtable discussion in May 2026. These perspectives have been synthesised for the purposes of this report. Unless expressly attributed, the views expressed in this report should not be attributed to any individual contributor, and the inclusion of any contributor does not imply endorsement of the report, its contents or conclusions.