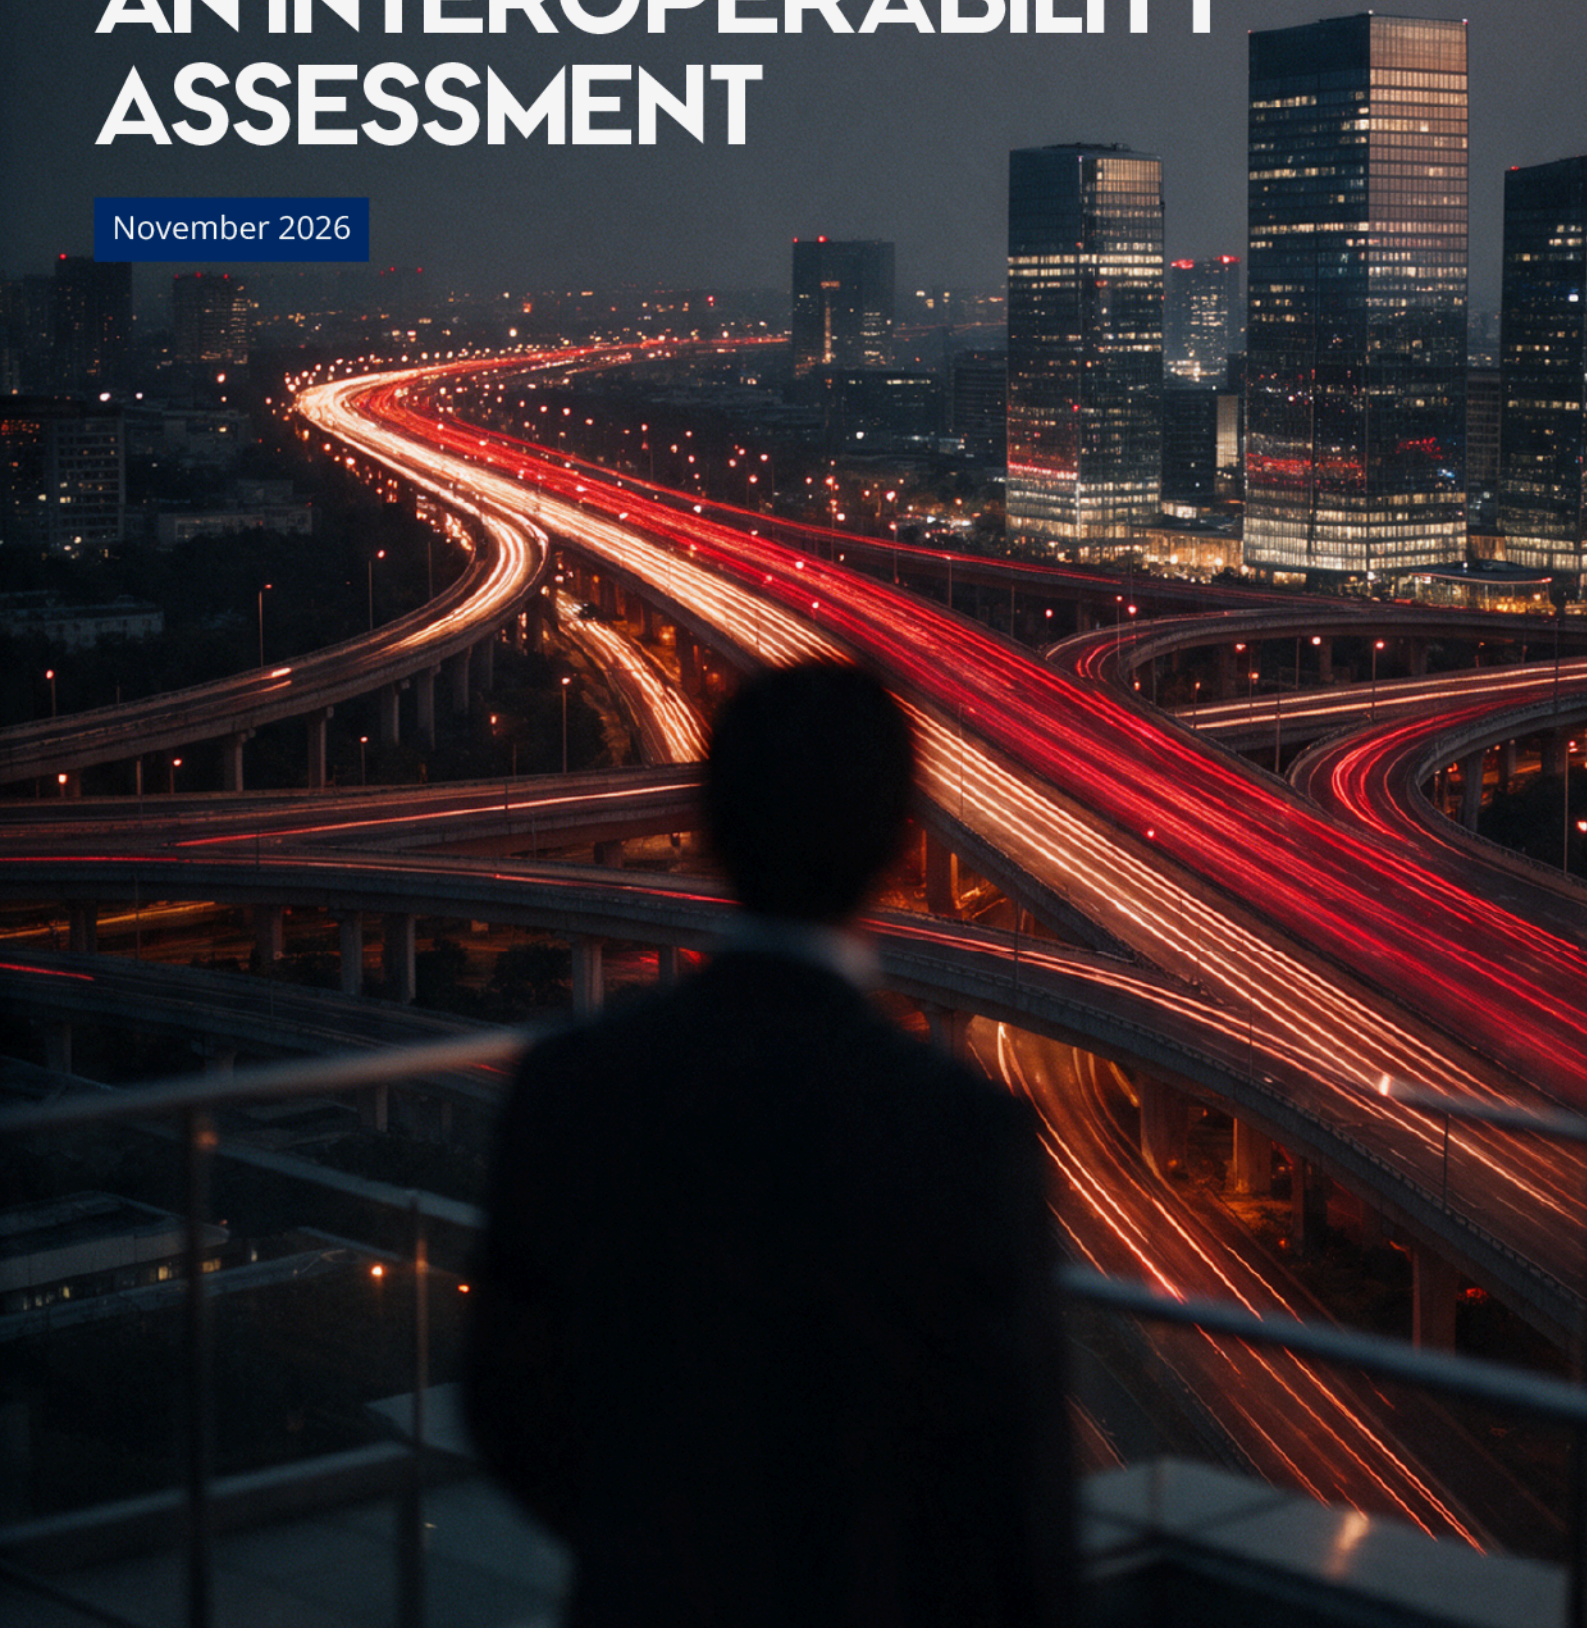
ASEAN Bankers Association

GFTN
Innovating Finance. For Everyone.

PHASE 1 REPORT

ASEAN AI GOVERNANCE: AN INTEROPERABILITY ASSESSMENT

November 2026



CONTENTS

1. Executive summary **03**
2. Why this matters: the cost of fragmentation and the value of interoperability **04**
3. Approach and methodology **06**
4. The state of play: how ASEAN markets govern AI **08**
5. The regional picture, pillar by pillar **09**
6. Case Study: A regional platform in practice **20**
7. Interoperability opportunities: Common standards, shared tools and supervisory alignment **21**
8. The convergence already under way **26**
9. Next steps **28**
10. Country profiles **28**
11. Annex A: Where regional capacity building would help most **29**
12. Annex B: How the opportunities trace to the evidence **30**
13. Sources and evidence notes **32**
14. Glossary **32**

© 2026 GFTN Limited, All Rights Reserved. Reproduction Prohibited.

01
EXECUTIVE SUMMARY

Artificial intelligence is moving to the centre of banking in Southeast Asia. It is already embedded across core banking operations, and generative and, increasingly, agentic systems are extending its reach into more of the work banks do. For a region whose financial sector sits within one of the world's fastest-growing digital economies, the opportunity is large, and adoption is already well advanced: around three-quarters of banks run generative AI in production or are piloting it.

How that AI is governed, though, is settled market by market. As banks increasingly build on the same models, cloud platforms and providers, and move data across borders to run them, whether the region's rules work together or pull apart becomes central to whether the sector can realise what AI offers. This report assesses how far AI governance across the eleven ASEAN markets is interoperable today. It follows the recommendations presented to the ASEAN Central Bank Governors and Financial Institutions CEO Dialogue in April 2026. These called on ASEAN banks, with support from national regulators, to work towards common principles and standards for AI governance in financial services.[1] The report draws on three sources. The first is a market-by-market reading of primary law and regulation. The second is a survey of senior technology, data and risk leaders at 61 banks across all eleven markets. The third is a set of interviews with supervisors in the region, held under the Chatham House Rule. Throughout, the report describes how the markets' rules converge or differ across borders. It does not rank or grade them.

The central finding is that the region holds more of the common ground interoperability requires than its separate rulebooks suggest, and that the differences that remain are few and clearly located. Five points carry the assessment.

Governing AI differently in each market carries a real cost for the banks that operate across the region. When the shared foundations of AI are governed differently from one market to the next, a bank meets the same underlying expectation in several different forms, which duplicates effort, constrains how data can be used, and slows deployment. The banks place a high value on AI being governed consistently across the markets they serve; it is a strong signal in the survey.[2] Interoperability is what would let the sector capture the value of AI across the region.

Across most of the eight pillars examined, the region already shares the ground a common approach would need. On governance and accountability, cyber security, human oversight and the expectation that AI be fair, the markets' requirements point broadly the same way, often because they draw on the same international standards and the same long-established banking law. Common ground also includes markets that have not yet written an AI-specific rule on a point, so it describes compatibility rather than maturity. Where a market differs, the difference is usually an overlay that a regional approach can absorb as an extra step, or a gap the whole region shares and could close together.

This common ground is the foundation for interoperability, and building on it will take deliberate effort. Shared requirements show how much room there is to align, but they do not let a bank rely across borders on work done once. Even where two markets ask for the same thing, each supervisor still assesses it against its own requirements, so a bank satisfies the same expectation separately in each market. Closing that gap needs regional arrangements that do not yet exist. One would let a check done in one market count in another, where supervisors choose to accept it. Another would give the markets shared definitions, so that they mean the same thing by the same words.

Only a few requirements call for a local implementation, and they concern data and the systems that hold it. Some markets require a bank's primary systems and data to be hosted onshore. In others, transferring customer data abroad is subject to the regulator's approval. A bank has to plan for the case where a transfer is not approved, so it needs storage and processing for that data in the country. Each market has its own reasons of security and sovereignty for these rules, and they place a real limit on how far a central set-up can reach. Banks feel the effect. Data rules are the largest source of cross-border friction they report, named by around two-thirds of them. Even so, a bank can keep its regional approach in these markets. It runs the affected systems and data processing locally, and keeps them aligned with its central set-up.

Much of the foundation is already in place, and the opportunities to build on it are beginning to emerge. Working independently, the markets have drawn on the same international references and joined the same regional instruments, so rules drafted separately have converged on a shared set of reference points. The report sets out three types of opportunity. Common standards are shared reference points that banks and supervisors could agree to use. Shared tools and skills are practical resources they could use together. Both could start in the near term. Supervisory alignment needs deeper agreement and trust between supervisors, and would take longer. The timing matters, because attention is already shifting to more autonomous, agentic AI, which concentrates the same risks the report traces. Much of what is needed to close the remaining gaps is already taking shape.

1 ASEAN Bankers Association, Development of an ASEAN Baseline Framework for Navigating AI Governance in Financial Services, discussion note presented to the ASEAN Central Bank Governors and Financial Institutions CEO Dialogue, 9 April 2026

2 The number of responses varies by market, but the result holds when each market is weighted equally. Responses are self-reported and are not a statistically representative sample (see Section 13).

© 2026 GFTN Limited, All Rights Reserved. Reproduction Prohibited.

02

WHY THIS MATTERS

The cost of fragmentation and the value of interoperability

Artificial intelligence is moving from the edges of banking to its core. It is already at work in credit decisions, fraud and financial-crime detection and customer channels, and generative and agentic systems are now extending that reach into document-heavy and decision-heavy work. The prize is large. McKinsey estimates that generative AI alone could add between US$200 billion and US$340 billion a year to the global banking sector, equivalent to between 9% and 15% of the industry's operating profits.[3] Southeast Asia is moving fast. According to the e-Conomy SEA 2025 report by Google, Temasek and Bain, the region's digital economy is on track to pass US$300 billion in gross merchandise value in 2025, growing at around 15% a year. AI accounted for over 30% of private funding value in the first half of 2025, and digital financial services continue to attract the largest share of private funding overall.[4]

This adoption is real, and it is deliberate. The Financial Stability Board (FSB) finds that banks are using AI mainly to strengthen internal operations and compliance, and are taking a measured approach to generative AI even as interest climbs.[5] What matters for governance is that this progress rests on shared foundations. Banks increasingly build on the same handful of models, cloud platforms and specialist vendors, and they move data across borders to train and run them. The Board identifies the systemic pressure points precisely there: third-party dependencies and the concentration of a few service providers, model risk and data governance, cyber exposure, and the risk of correlated behaviour when many institutions rely on the same tools.

Its October 2025 monitoring report returns to the same concern as a small number of global technology providers come to control more of the AI supply chain.[6] The International Monetary Fund reached a similar conclusion in its October 2024 Global Financial Stability Report, flagging a reliance on a few key third-party AI providers and the risk that rapid, correlated AI-driven decisions produce herd-like behaviour.[7] Supervisory responses are still at an early stage. The Bank for International Settlements finds that most financial authorities have not yet issued AI rules specific to financial institutions.[8]

When each market governs those shared foundations differently, the friction lands on the banks that operate across them. The clearest evidence is in data. The Organisation for Economic Co-operation and Development (OECD) estimates that a fully fragmented regime for data could reduce global GDP by 4.5%, whereas open regimes backed by trusted safeguards could raise it by 1.7%.[9] The pattern holds within this region. The Asian Development Bank finds that stricter data-localisation and transfer rules raise operational costs and dampen innovation across Asia and the Pacific, and that the burden of fragmented compliance falls hardest on smaller firms.[10] For a bank running one AI model, or one fraud-detection system, across several ASEAN markets, divergent rules convert directly into duplicated compliance, constrained data and slower deployment.

The upside of alignment is just as concrete, and it is regional. ASEAN's Digital Economy Framework Agreement could double the region's digital economy to around US$2 trillion by 2030,[11] with trusted cross-border data flows among its core enablers.[12] Its negotiations concluded in May 2026, and it is due to be signed at the ASEAN Summit in November 2026. Banking sits at the centre of that digital economy. For banks, interoperable AI governance could mean that a model validated in one market can be trusted in another, and that data can move between markets under shared safeguards. That is what would let the sector capture the value on offer. Testing how far that interoperability exists today, market by market, is the task this report sets itself.

To ground the argument in what the region's banks are experiencing, the Global Finance & Technology Network (GFTN) and the ASEAN Bankers Association (ABA) surveyed senior technology, data and risk leaders at 61 banks across all eleven ASEAN markets. The survey finds that adoption is already well advanced. Around three-quarters of banks run generative AI in production or are piloting it, making it the most widely adopted of the seven uses the survey asked about. Customer-facing tools and operational uses such as know-your-customer checks follow, each at around half of banks (Exhibit 1).

This adoption has moved ahead of the rules meant to govern it, and the gap is already costing banks. Around half of those surveyed have had an AI deployment blocked, delayed or scaled back because of AI-related regulatory uncertainty or a difference between markets (Exhibit 2). Banks without cross-border operations report this cost too. Among them, 60% say an AI deployment was affected, against 47% of banks that operate across borders. For banks working in one market, the uncertainty they report can only concern that market's own rules. The banks are equally clear about what would help them most. Asked how much it matters that AI is governed consistently across the markets they operate in, they put the average at 4.3 out of 5, a strong signal. One bank noted that its AI use is still **"primarily limited to productivity and knowledge-management use cases"**, but that **"greater consistency would become**

3 McKinsey & Company, Capturing the full value of generative AI in banking, 2023
4 Google, Temasek & Bain, e-Conomy SEA 2025, 2025
5 Financial Stability Board, The Financial Stability Implications of Artificial Intelligence, 2024
6 Financial Stability Board, Monitoring Adoption of Artificial Intelligence and Related Vulnerabilities in the Financial Sector, 2025
7 International Monetary Fund, Global Financial Stability Report, October 2024
8 Bank for International Settlements, Regulating AI in the financial sector: recent developments and main challenges (FSI Insights No. 63), 2024
9 Organisation for Economic Co-operation and Development and World Trade Organization, Economic Implications of Data Regulation, 2025
10 Asian Development Bank, The Role and Future of Digital Economy Agreements in Developing Asia and the Pacific, 2025
11 ASEAN, ASEAN DEFA study projects digital economy leap to US$2tn by 2030, 2023
12 World Economic Forum, Why ASEAN's new Digital Economy Framework Agreement is a game-changer, 2025

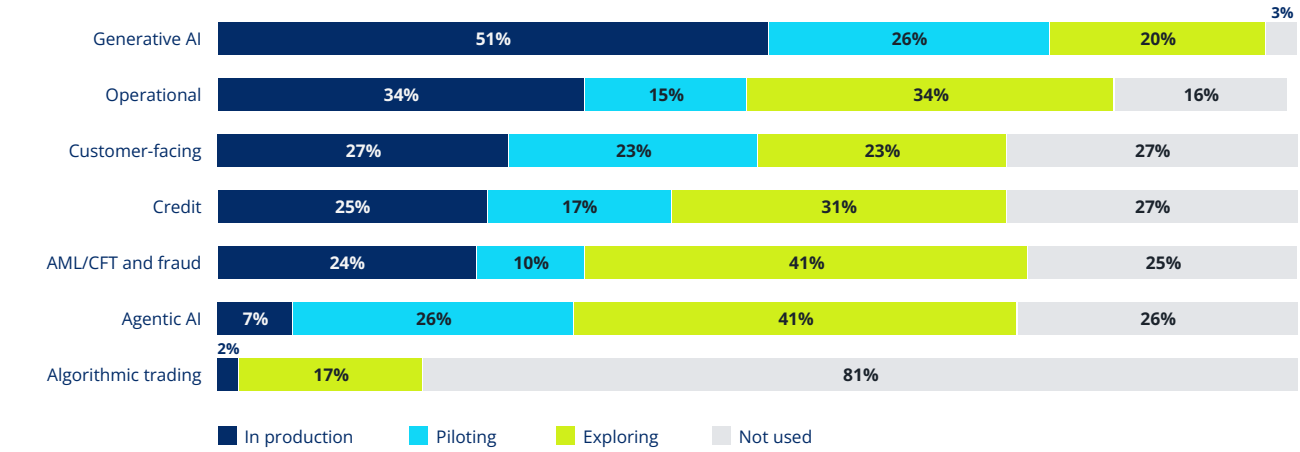
© 2026 GFTN Limited, All Rights Reserved. Reproduction Prohibited.

increasingly important if AI were expanded into customer-facing or risk-sensitive business processes."

Exhibit 1:

Stage of AI adoption, by use case

Share of the 61 surveyed banks at each stage, for each use case.

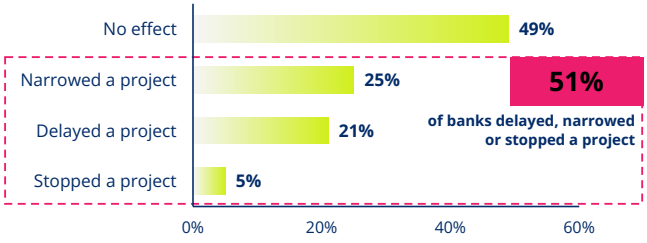


Source: the GFTN and ABA survey of ASEAN bank leaders.

Exhibit 2:

Whether AI-related regulatory uncertainty, or a difference between markets, has affected an AI deployment

Share of the 61 surveyed banks giving each answer; each bank gave one answer.



Source: the GFTN and ABA survey of ASEAN bank leaders.

The picture from the banks matches the case this section has set out. They have adopted AI across their operations, they carry a real cost when the rules are unclear or differ between markets, and they place a high value on governance moving closer together.

© 2026 GFTN Limited, All Rights Reserved. Reproduction Prohibited.

03

APPROACH AND METHODOLOGY

This report rests on a single comparative framework, applied in the same way to every market so that findings line up and interoperability can be read across borders. It combines the FSB's view of AI risk in finance with the principles the region has agreed for itself. It is used to describe where markets align and where they differ, without ranking them against one another.

THE EIGHT-PILLAR FRAMEWORK

The report takes the FSB's work on artificial intelligence in finance as its organising lens.[13] FSB examines AI through the concerns central to prudential supervision: the concentration of a small number of third-party providers, correlated behaviour across institutions, cyber exposure, and model risk and data governance. For its criteria, the report draws on the region's own reference point, the ASEAN Guide on AI Governance and Ethics of 2024 and its 2025 expansion for generative AI,[14,15] which set out seven guiding principles: transparency and explainability, fairness and equity, security and safety, human-centricity, privacy and data governance, accountability and integrity, and robustness and reliability.

Bringing these two reference points together produces the eight pillars that structure the report. Seven correspond to the ASEAN Guide's principles: governance and accountability, model risk, transparency, human oversight, fairness, data, and cyber security. The report adds an eighth dimension, third-party and concentration risk, to reflect the supervisory significance of banks' growing reliance on shared cloud and model providers. Together, the eight pillars cover the principles the region has agreed to and the concerns central to banking supervision. The report applies them to every market, so that the markets can be compared on the same terms.

MAPPING THE PILLARS IN EACH MARKET

Against each pillar, the report describes how a market's rules relate to a cross-border approach, using three descriptions of fit. The report looks at interoperability from the position of a bank serving several ASEAN markets from a central regional set-up. That set-up can be a technology platform, but it also includes the bank's central processes and controls.

Common ground. The central set-up can serve the market under its rules as they stand. This includes requirements that every market shares. It also includes markets that have not yet written a rule for AI on a given point. Where the region shares such a gap, the report says so in the pillar concerned. Even on common ground, each supervisor still assesses compliance separately, so full interoperability would need further steps, such as mutual recognition.

Overlay. A market-specific requirement that can be met within the central set-up through an additional step, such as a filing, a disclosure or a reporting line, without a separate in-country implementation.

Local implementation. A market-specific requirement that cannot be met within the central set-up, even with additional steps, because part of it has to be established and run separately in the country. Hosting primary systems onshore is an example.

An overlay adds something to the central set-up, whereas a local implementation replaces part of it with a local equivalent. These are descriptions of how each market's framework fits with those of its neighbours, rather than an assessment of quality. The categories, and the placement of each market within them, reflect the report's interpretation of each market's instruments. They are based on the rules in force at the time of writing, and several of those rules are still evolving.

THE COUNTRY PROFILES

The report draws on eleven country profiles, one for each ASEAN market, each built on the same eight-pillar structure. Within a profile, a market's dedicated AI-governance document sets out how its regulator approaches AI, while the binding instruments beneath it determine what institutions are required to do. These include banking laws, central-bank notices and circulars, data-protection laws and outsourcing rules. The interoperability findings follow these binding instruments where they exist, because they are what a bank must satisfy in each market and so are the main source of cross-border friction. Where a market sets its expectations in supervisory guidance instead, the findings follow that guidance. Where a market has no dedicated AI-in-banking instrument, the report notes this, as the absence is itself a finding.

The profiles rely on official and primary sources: the government and regulatory texts themselves. Each instrument is dated and checked for whether it remains in force, and each finding reflects how firmly it is established, from binding law through to published guidance and observed practice. Sources in languages other than English were consulted in the original.[16]

PRIMARY RESEARCH: BANK LEADERS AND REGULATORS

Alongside the desk research, the report draws on two strands of primary input, chosen to capture both sides of the picture. The

13 Financial Stability Board, The Financial Stability Implications of Artificial Intelligence, 2024
14 ASEAN, ASEAN Guide on AI Governance and Ethics, 2024
15 ASEAN, Expanded ASEAN Guide on AI Governance and Ethics (Generative AI), 2025
16 Where an official text was not publicly available, the report relies on reputable secondary legal sources and cites them as such.

© 2026 GFTN Limited, All Rights Reserved. Reproduction Prohibited.

first is a survey of bank leaders. GFTN and the ASEAN Bankers Association surveyed senior technology, data and risk leaders at 61 banks across all eleven ASEAN markets, asking how far they have taken AI, where they find the rules clear or unclear, and what the differences between markets cost them in practice. The second is a set of interviews with the region's supervisors. GFTN spoke with central banks and financial regulators in several markets to hear the official view on scope, legal status and oversight. These interviews were held under the Chatham House Rule, so the report draws on what was said without attributing it to any individual or institution.

These inputs are complemented by a series of GFTN Select roundtables across the region, each including a dedicated session on AI governance.

© 2026 GFTN Limited, All Rights Reserved. Reproduction Prohibited.

04

THE STATE OF PLAY

How ASEAN markets govern AI

Before turning to the individual pillars, here is the broad shape of how the eleven markets govern artificial intelligence in banking today: the institutions responsible for it and the different approaches they take. The markets differ in their sizes, priorities and points of departure, and the intention here is to describe that landscape rather than to measure one market against another. The specific instruments in each market, and their sources, are set out pillar by pillar in the sections that follow, and in the country profiles.

WHO HOLDS RESPONSIBILITY

In most ASEAN markets the central bank or financial regulator is the primary authority for AI in banking. Examples include the Monetary Authority of Singapore, Bank Negara Malaysia, the State Bank of Vietnam, Bangko Sentral ng Pilipinas, the National Bank of Cambodia and the Banco Central de Timor-Leste. The national data-protection authority typically shares responsibility for data and privacy. Beyond that common pattern, the division of responsibility varies. In some markets a single authority holds almost the entire agenda, addressing AI within its own technology and risk framework, while in others it is shared across several bodies. In a few, AI-specific measures come from outside the financial regulator, for example from a ministry responsible for technology. In Brunei, a distinct framework for Islamic banking applies alongside the conventional rules. These are differences of institutional design, each reflecting a market's own structure and priorities.

DIFFERENT APPROACHES TO REGULATION

The markets also differ in how they regulate, along a familiar spectrum. Some set out their expectations mainly through guidance and principles, leaving institutions room to interpret and apply them, while others favour more prescriptive, binding requirements. A number have introduced measures written specifically for AI. Some have practical frameworks and methods for testing and assurance. One has a binding prudential requirement on how banks deploy AI, and another has defined AI in law. More commonly, AI is addressed through the frameworks that already govern banking, such as banking acts, risk-management and outsourcing rules, data-protection laws and cyber requirements. This shared reliance on established law is itself something the region has in common. The frameworks also differ in their reach. Some extend to digital and challenger banks, financial-technology firms and the local operations of foreign banks, and others do not. The pillars below take up this question where it matters.

Where a market has kept its expectations at the level of guidance, that is often a deliberate choice. One supervisor described holding its AI guidance to principles so as not to become **"the regulatory burden that would make banks stop using AI."** The banks work within that space. Where the rules are still forming, they tend to proceed through controlled, low-risk pilots and to lean on international frameworks until clearer expectations arrive. Adoption and policy tend to go together. In the bankers' survey, 86% of banks with five or more AI use cases live or in pilot have a governance policy in force in at least one area, against 31% of banks with two or fewer.

THE OVERALL PICTURE

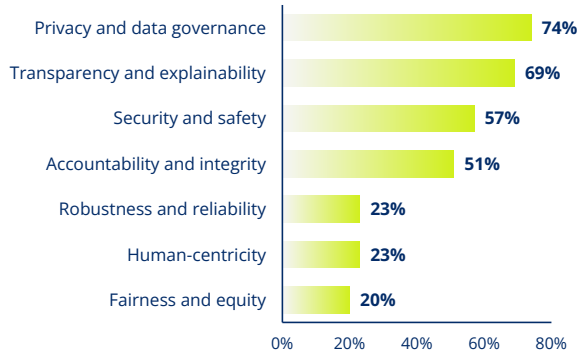
Taken together, a consistent pattern emerges. Across most of the region there is already substantial common ground on the fundamentals of sound governance: on accountability for what AI does, on cyber security, and on the expectation that AI should be fair. The picture on model risk, the discipline of testing and validating a model before relying on it, is still taking shape, with some markets setting out detailed expectations and others relying on their general risk rules. The requirements that call for a local implementation concern where a bank's systems and data must be hosted, and whether customer information can move across borders. On these questions markets range from the openly accommodating to those that set more specific conditions. The section that follows examines each of these pillars in turn.

The banks are equally clear about where they most want the rules firmed up. Privacy and data governance comes first, then transparency and explainability, with the security and safety of AI systems and accountability close behind (Exhibit 3). Their top two priorities are also where national rules vary most: eight of the eleven markets set their own requirements on data governance or on where data and systems must be located, and seven do so on transparency (Exhibit 6).

Exhibit 3:

Areas that would benefit most from greater regulatory clarity

Share of the 61 surveyed banks that named each area; each bank could choose up to three.



Source: the GFTN and ABA survey of ASEAN bank leaders.

© 2026 GFTN Limited, All Rights Reserved. Reproduction Prohibited.

05

THE REGIONAL PICTURE, PILLAR BY PILLAR

Across the eight pillars, one pattern runs through the detail that follows. On most of them the eleven markets already share the ground a common approach would need, and where a market differs it is usually an overlay that the central set-up can absorb, or a gap the whole region shares and could close together. Requirements that call for a local implementation are rare. Some markets require a bank's primary systems and data to be hosted in the country. In others, transferring customer data abroad is subject to the regulator's approval. A bank has to plan for the case where a transfer is not approved, so it needs storage and processing for that data in the country, established and run locally. Shared ground of this kind shows how much room there is to align. It does not by itself make the rules work together: one regulator does not automatically accept what another has approved, and even identical requirements can oblige a regional bank to satisfy them market by market. The pillar-by-pillar account below is the evidence for that pattern, and Exhibit 6 at the close of the section draws it together.

Each pillar below follows the same structure. It opens with what the pillar means and why it matters. It then describes the common ground across the region, the overlays some markets add, and any requirements that call for a local implementation. Exhibit 4 summarises the eight pillars.

Exhibit 4:

The eight pillars at a glance

Pillar	In one line
Governance and accountability	Clear responsibility, at board and senior-management level, for how AI is used
Model risk	Testing and validating an AI model before relying on it, and monitoring it in use
Transparency	Being able to explain what an AI system does, and disclosing when it is used
Human oversight	Keeping a person able to review, intervene in, or override an AI-driven decision
Fairness	Guarding against AI that produces biased or discriminatory outcomes
Data	How customer information is governed and protected, including whether it can move across borders
Cyber security	Protecting AI systems and the data behind them from attack and compromise
Third-party and concentration risk	Managing reliance on a small number of shared cloud and AI providers

Source: GFTN analysis of the eleven country profiles.

GOVERNANCE AND ACCOUNTABILITY

Governance and accountability is the requirement that clear responsibility for how AI is used sits with a bank's board and senior management. Responsibility for what an AI system does cannot be transferred to the system itself: when an automated decision causes harm, someone must still answer for it. The FSB's 2026 sound practices for the responsible adoption of AI, finalised in October 2026, are addressed in the first instance to the board and senior management of financial institutions, tying the safe use of AI to the oversight exercised at the top of the firm.[17]

This is the pillar on which the region has the most in common. Every market already fixes board and senior-management responsibility through the conventional corporate-governance and fit-and-proper requirements long embedded in banking law, modelled in several markets on the international principles of the Basel Committee and the OECD. These requirements set the familiar structure that AI accountability attaches to. The board must be actively engaged. Directors and senior officers must pass the regulator's fit-and-proper test of honesty, competence and financial soundness before they are approved. Most markets require board audit and risk committees. The rest expect them, or give the same oversight to an equivalent body. The three lines of defence give distinct responsibilities to the business that takes on

17 Financial Stability Board, Sound Practices for Responsible Adoption of Artificial Intelligence (AI): Final report, 2026

© 2026 GFTN Limited, All Rights Reserved. Reproduction Prohibited.

risk, to the risk and compliance functions that oversee it, and to an independent internal audit. Cambodia's governance and internal-control rules,[18] Thailand's binding corporate-governance framework resting on its Financial Institutions Business Act,[19] and Indonesia's governance regulation for commercial banks[20] are representative, and Timor-Leste's 2025 update of its corporate-governance instruction shows the same structure being renewed.[21] Because outsourcing a function does not remove responsibility for it, a bank cannot contract its accountability away by handing model development to a vendor, a principle several markets state expressly.

Several markets go beyond this, and they have done so in two broad ways. The first is to pin responsibility on an identified person. Singapore and Malaysia both have individual-accountability frameworks that provide a foundation for allocating responsibility for AI-related activities, although neither was designed specifically for AI. Singapore's guidelines on individual accountability expect a bank, on a comply-or-explain basis, to assign each core function to a named senior manager and to hold that manager accountable for it. These include the chief risk, technology and data officers, who oversee model, technology and data governance.[22] Singapore's 2026 guidelines on AI risk management supplement this foundation with direct expectations for board and senior-management oversight of AI risk, including clear roles, responsibilities and risk-management arrangements.[23] Malaysia's responsibility-mapping requirements are binding. They oblige the board to allocate every material responsibility to a named, fit-and-proper officer.[24] Malaysia's technology-risk rules also give the board and senior management specific duties, including a designated chief information security officer. These duties cover technology and cyber risk generally, including the risks of new or emerging technology.[25] In both, accountability is personal and cannot be delegated to the machine. The second approach is to fold AI into the board oversight that already exists. Indonesia's AI governance handbook, which is guidance, maps AI supervision onto the board committees that banks must already have, so that AI stays a standing item on the board's agenda.[26] Thailand's advisory AI guideline asks the board and senior management to own AI risk within the firm's risk appetite.[27] The Philippine central bank's voluntary governance principles for AI, issued in June 2026, likewise hold management ultimately accountable for decisions informed by AI, and recommend that each bank set up an AI governance framework proportionate to its use of AI.[28] Vietnam has put the allocation into statute, its Law on Artificial Intelligence assigning duties along a chain of developer, provider, deployer and user, a binding division of responsibility for AI that is rare in the region.[29]

The local variations add to this shared foundation. Vietnam's statute also puts the local entity that deploys a high-risk system first in line to answer if it goes wrong. A regional group running a shared model would therefore agree internally that the local entity can recover from whichever group company built the model. Several markets, among them the Philippines, treat a regionally located decision-maker, a country head or an officer of equivalent rank, as the functional board of a locally operating branch, which allows a group operating model to satisfy the local accountability requirement without relocating its people.[30] In the markets whose banking law does not yet name AI, accountability rests by default on these same general governance duties, which reach an AI-driven decision as they would any other. Because every market shares this corporate-governance foundation, and because the markets that have named AI have anchored their expectations to it, governance and accountability is the pillar on which a common regional reference point would meet the least resistance.

Cross-cutting theme

ISLAMIC BANKING: A SHARED SHARIAH-GOVERNANCE OVERLAY

In the markets where Islamic banking is significant, an Islamic bank carries an additional, binding layer of governance above the conventional one, and this layer is notably consistent in shape across the region. Malaysia's Shariah Governance policy document requires Shariah compliance to run through the whole institution, including its operations, risk management and internal controls as well as its products. A Shariah Committee is the binding religious authority. Dedicated Shariah review, audit and risk-management functions sit alongside the conventional ones, and the board is ultimately accountable.[31] Brunei applies the same design through its Syariah Governance Framework, with a Syariah advisory body and Syariah control functions whose accountability rests with the board and a statutory Syariah authority as the final arbiter.[32] Indonesia's 2024 sharia-governance regulation sets a parallel structure for its Islamic banks and business units, built around each institution's Sharia Supervisory Board and the same review, audit and risk functions.[33]

For AI, the implication is the same in each. A model an Islamic bank uses to price, score or structure a product must operate within Shariah-compliant parameters and falls within the reach of Shariah oversight and audit, so it has to be explainable and auditable in Shariah terms, and its accountability cannot pass to the model or to a vendor. The structure is shared, but the binding rulings are national: each market's Shariah authority has the final word, and these determinations are not the subject of regional mutual recognition. A regional bank cannot therefore rely on a home-market Shariah sign-off for its Islamic-banking activity in another market, and a shared AI system serving that activity has to leave room for local Shariah oversight, audit and parameters on top of its conventional controls. This is a manageable overlay, and its common shape across markets makes it a natural candidate for the kind of mutual recognition that does not yet exist.

18 National Bank of Cambodia, Prakas on Governance in Banks and Financial Institutions, 2008
19 Bank of Thailand, Notification FPG 1/2566 on Corporate Governance of Financial Institutions, 2023
20 Otoritas Jasa Keuangan, Regulation No. 17 of 2023 on the Implementation of Governance for Commercial Banks, 2023
21 Banco Central de Timor-Leste, Instruction No. 29/2025 on the Corporate Governance of Banks, 2025
22 Monetary Authority of Singapore, Guidelines on Individual Accountability and Conduct, 2020
23 Monetary Authority of Singapore, Guidelines on Artificial Intelligence Risk Management, 2026
24 Bank Negara Malaysia, Policy Document on Responsibility Mapping, 2023
25 Bank Negara Malaysia, Risk Management in Technology (RMiT), policy document, 2026
26 Otoritas Jasa Keuangan, Artificial Intelligence Governance for Indonesian Banking, 2025
27 Bank of Thailand, Policy Guideline on Risk Management of the Use of AI Systems, 2025
28 Bangko Sentral ng Pilipinas, Memorandum No. M-2026-031, Governance Principles for Artificial Intelligence in Financial Services, 2026
29 Vietnam, Law on Artificial Intelligence No. 134/2025/QH15, 2025
30 Bangko Sentral ng Pilipinas, Manual of Regulations for Banks, Section 148 (IT Risk Management, consolidating Circular No. 1137), 2022
31 Bank Negara Malaysia, Shariah Governance, policy document BNM/RH/PD 028-100, 2019
32 Brunei Darussalam Central Bank, Notice on the Syariah Governance Framework (IFAU/N/1/2018, Amendment No. 1), 2024
33 Otoritas Jasa Keuangan, Regulation No. 2 of 2024 on the Implementation of Sharia Governance for Sharia Commercial Banks and Sharia Business Units, 2024

© 2026 GFTN Limited, All Rights Reserved. Reproduction Prohibited.

MODEL RISK

Model risk is the discipline of testing and validating an AI model before relying on it, and keeping it under review once it is in use. A model can look authoritative and still be wrong, and organisations are already feeling the consequences. McKinsey's 2025 global AI survey found that just over half of respondents from organisations using AI had seen at least one negative consequence from it, and nearly a third reported consequences stemming from AI inaccuracy.[34] An AI model used to price credit, screen for fraud or value assets therefore has to be tested and monitored rather than simply trusted.

On this pillar the requirements are uneven. A few markets set out detailed duties to validate AI models, while most rely on their general risk-management rules, so the approaches do not conflict even though they differ in depth. The common starting point is a general risk-management or technology-risk framework that treats an AI model like any other system, with no specific duty to validate the model itself. The region does not yet share a baseline written for AI models.

Vietnam has introduced a binding model-risk regime, rare in the region. The State Bank of Vietnam's internal-control circular requires a bank to keep a tiered inventory of its models and to report on model risk. Each high-risk model must be validated before use and at least once a year, by a second line independent of the people who built it.[35] Singapore,[36] Thailand[37] and Cambodia[38] address model risk through guidance. Singapore's guidelines, issued in October 2026, set supervisory expectations across the AI life cycle and take effect in stages from October 2027. Brunei takes a distinctive route of its own, a binding requirement that an AI system on a critical function demonstrate low error rates over a sustained period before it may run more autonomously.[39] Timor-Leste applies a related discipline in a narrower setting, requiring independent validation of its credit and accounting models.[40]

In several other markets, work on model risk is under way. Malaysia's central bank has published a discussion paper on artificial intelligence in the financial sector.[41] The paper shares the findings of an earlier survey by the central bank and asks for public views on the proposed approach. The Philippine central bank consulted in 2025 on a draft circular on model-risk management, which has not yet been finalised.[42] Indonesia's regulator has released a handbook on AI governance for banking.[43] All three draw on the same international references. Whether binding rules follow is not yet settled. For now, banks in these markets rely on general risk-management frameworks. Separately, and beyond the banking rulebook, Malaysia consulted in July 2026 on a proposed cross-sectoral AI Governance Bill, which would move the country from voluntary guidelines towards binding obligations across the AI life cycle.[44] The government expects to table it in Parliament by early 2027.[45] Thailand's Electronic Transactions Development Agency published a draft cross-sectoral AI Act for public hearing in July 2026, which would set risk-based obligations, including prohibited and high-risk categories. It has not yet been enacted.[46]

Because model risk is a gap the region largely shares rather than a point of conflict, and because several of the emerging rules are built to recognise a group or parent framework, it is one of the more promising pillars for a common reference point.

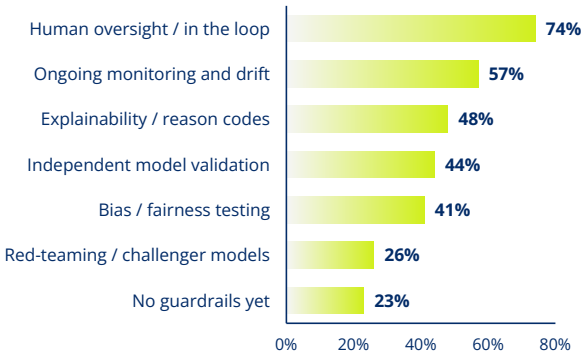
The banks' own practice shows the same gap from the inside. Around three-quarters of them apply human oversight to higher-risk AI, while the more formal model controls are further behind: on the survey, fewer than half apply independent model validation to higher-risk AI and under a third apply red-teaming or challenger models (Exhibit 5). One bank said the wider roll-out of AI in its customer-facing and credit work had been **"paced/delayed pending clarity"** on explainability, independent model validation and permitted data use. A shared reference on model risk would give a bank like this a firmer footing to build on.

Formal testing is far more common among banks with a high-stakes use case or a policy in place. Among the 26 banks with AI credit models live or in pilot, 73% validate their models and 73% test for bias. Among the other 35 banks, 23% validate and 17% test for bias. Of banks with a fairness policy in force, 74% test for bias, against 7% of banks with none. These patterns show association, not cause. Testing is thinnest among banks with neither. Supervisors hear the same need. One said banks are asking **"how can institutions determine the materiality and risk level of AI use cases"** and **"what kind of standards do we expect for model validation, model monitoring"**.

Exhibit 5:

Guardrails applied to higher-risk AI

Share of the 61 surveyed banks that apply each guardrail; banks could select all that apply.



Source: the GFTN and ABA survey of ASEAN bank leaders.

34 McKinsey & Company, The state of AI in 2025: Agents, innovation, and transformation, 2025
35 State Bank of Vietnam, Circular 83/2025/TT-NHNN on internal control systems, 2025
36 Monetary Authority of Singapore, Guidelines on Artificial Intelligence Risk Management, 2026
37 Bank of Thailand, Policy Guideline on Risk Management of the Use of AI Systems, 2025
38 National Bank of Cambodia, Technology and Cyber Risk Management Guidelines, 2026
39 Brunei Darussalam Central Bank, Notice on Technology Risk Management (TRS/N-1/2023/2), 2023
40 Banco Central de Timor-Leste, Instruction No. 16/2021 on Credit Risk Management, 2021
41 Bank Negara Malaysia, Discussion Paper on Artificial Intelligence in the Malaysian Financial Sector, 2025
42 Bangko Sentral ng Pilipinas, Draft Circular on Model Risk Management Guidelines, 2025
43 Otoritas Jasa Keuangan, Artificial Intelligence Governance for Indonesian Banking, 2025
44 Ministry of Digital (Malaysia), National AI Office, Public Consultation on the Proposed Artificial Intelligence (AI) Governance Bill, 2026
45 Malay Mail, Digital minister says Malaysia's draft AI Governance Bill complete, hopes to table it by Q1 2027, 2026
46 Baker McKenzie, Thailand: Draft AI Act, Seven Key Implications (on the Electronic Transactions Development Agency draft AI Act, public hearing to 14 August 2026), 2026

© 2026 GFTN Limited, All Rights Reserved. Reproduction Prohibited.

TRANSPARENCY

Transparency is the ability to explain what an AI system does and to disclose when it is being used. A decision that cannot be explained cannot be questioned: a customer refused credit by an opaque model has no way to understand or challenge the outcome, and the bank has no way to give the reasons the law requires. In 2023 the United States Consumer Financial Protection Bureau set out that a lender using AI or a complex algorithm must still give a customer the specific and accurate reasons for a credit denial, with no special exemption for artificial intelligence.[47] A model a bank relies on has to be explainable enough to justify its decisions, however sophisticated it is.

Every market starts from the same conventional foundation: binding rules requiring clear, fair and non-misleading disclosure of a product's terms, fees and risks, so that customers are told what they are buying. These product-and-conduct disclosure regimes are aligned with international practice. Brunei's rules, for example, set minimum pre-contractual disclosures and call for a product disclosure sheet.[48] What these rules do not address is the explanation of an AI-driven decision, and the markets take different approaches to it.

Where an AI-specific transparency duty exists, it shares two features across the region: it usually takes the form of guidance instead of binding law, and the duty tends to fall to the data-protection regulator rather than the banking regulator. In the Philippines, the privacy regulator expects a bank using AI on personal data to explain, in plain language, the factors the system weighs, which demands real insight into how the model works. Its data-protection rules go further, requiring the logic of automated processing to be filed with the regulator in advance.[49,50] Singapore takes a proportionate approach: its data-protection regulator encourages firms to disclose when AI is used,[51] and its AI risk management guidelines calibrate the degree of transparency and explainability required to the context and nature of the AI use case and its risk materiality, so that AI use with higher impact on customer outcomes such as credit decisioning carries a higher bar.[52]

Malaysia's 2026 automated-decision-making guideline likewise expects a bank to tell a customer when a decision involves automated processing or profiling.[53] In each, the duty is owed through the privacy channel. Indonesia and Thailand set the expectation through the banking regulator instead. Indonesia's AI governance handbook makes transparency and explainability core expectations for a bank's AI systems,[54] and Thailand's AI policy guideline asks a bank to be able to explain its models' outputs and to test and monitor them for reliability.[55] Both are guidance rather than binding requirements.

A newer strand of transparency duty is starting to appear, aimed at generative and conversational AI. Cambodia's 2026 technology guidelines require a bank to tell customers when they are dealing with a chatbot instead of a person,[56] and Vietnam's 2025 Law on Artificial Intelligence requires a system that interacts with people to make clear that it is a machine and requires AI-generated audio, image and video to be marked.[57] Vietnam separates protecting a model from explaining its decisions. Its AI law lets a firm keep its source code, detailed algorithms and training data confidential. A duty to explain an automated decision to the customer still applies, but it sits in Vietnam's personal-data rules and is supervised by the security ministry rather than the banking regulator, so a bank meets it through its privacy function.[58]

The conventional disclosure regime is shared across the region and interoperable. The AI-specific transparency requirements are still limited and largely expressed as guidance, and because they usually arise under privacy law, a common explainability standard set by banking regulators alone would not meet them; the data-protection authorities would need to be involved as well. For a bank operating across several markets, the requirement is manageable: it can be met by extending existing disclosure and model-documentation practice, and by handling the explanation of automated decisions through the privacy function where local law directs.

HUMAN OVERSIGHT

Human oversight is the requirement to keep a person able to review, intervene in or override an AI-driven decision. A decision taken entirely by a machine, with no one able to step in, leaves an error uncorrected and the customer without recourse. This is why regulators are placing a person back into consequential decisions: the European Union's 2024 AI Act treats AI used to score a person's creditworthiness as high-risk and requires such systems to be built so that a person can oversee them and intervene.[59] How much oversight a decision needs rises with what is at stake for the customer: a routine, high-volume decision can run with a light touch, while a decision that materially affects someone calls for a competent person able to check it and, if needed, reverse it.

Every market gives a customer a way to challenge a decision. Each has a binding complaints and redress regime, so a customer affected by an AI-influenced decision can dispute the outcome, ask for it to be reviewed and, if still dissatisfied, escalate to the regulator or a dispute-resolution body. Brunei's complaints rules require independent handling, written reasons for a refusal and a right to escalate to the central bank,[60] and Cambodia's set a fixed resolution timeline and a route on to the regulator and the courts.[61]

47 US Consumer Financial Protection Bureau, Guidance on Credit Denials by Lenders Using Artificial Intelligence, 2023
48 Brunei Darussalam Central Bank, Guidelines on Product Transparency and Disclosure (FCI/G2/2019/1), 2019
49 National Privacy Commission, Advisory No. 2024-04, Guidelines on Artificial Intelligence, 2024
50 Philippines, Data Privacy Act of 2012 (RA 10173) and Implementing Rules and Regulations, 2016
51 Personal Data Protection Commission (Singapore), Advisory Guidelines on Use of Personal Data in AI Recommendation and Decision Systems, 2024
52 Monetary Authority of Singapore, Guidelines on Artificial Intelligence Risk Management, 2026
53 Personal Data Protection Commissioner (Malaysia), Automated Decision-Making and Profiling Guideline, 2026
54 Otoritas Jasa Keuangan, Artificial Intelligence Governance for Indonesian Banking, 2025
55 Bank of Thailand, Policy Guideline on Risk Management of the Use of AI Systems, 2025
56 National Bank of Cambodia, Technology and Cyber Risk Management Guidelines, 2026
57 Vietnam, Law on Artificial Intelligence No. 134/2025/QH15, 2025
58 Vietnam, Decree 356/2025/ND-CP implementing the Law on Personal Data Protection, 2025
59 European Union, Regulation (EU) 2024/1689 (Artificial Intelligence Act), 2024
60 Brunei Darussalam Central Bank, Notice for the Establishment of a Complaints Handling Function (FCI/N1/2021/1), 2021
61 National Bank of Cambodia, Prakas on Resolution of Consumer Complaints, 2017

© 2026 GFTN Limited, All Rights Reserved. Reproduction Prohibited.

Beyond this shared baseline, the markets that address the automated decision itself take two approaches. Some constrain the decision directly. In the Philippines, a fully automated decision with legal effect for a customer has been barred since 2016 unless the customer consents.[62] The privacy regulator adds that human intervention must be meaningful and carried out by people with the necessary competence and authority, and that customers must be able to question and contest automated decisions that pose a significant risk to them.[63] Vietnam frames human control as a founding principle of its 2025 AI Law, requiring that an authorised person be able to review, intervene in, refuse or alter a system's decision before it takes effect.[64]

Others require attention to automated decisions without prohibiting them. Indonesia gives a customer the right to object to a decision taken solely by automated means where it produces a legal or similarly significant effect, such as a refusal of credit.[65] Malaysia's 2026 guideline applies a broad trigger to automated decisions while treating human review as expected good practice.[66] Singapore expects appropriate human oversight that takes into account the objective of using AI, and is proportionate to the risk of each AI use case.[67] Thailand leaves meaningful human involvement to be considered case by case according to risk.[68]

These approaches converge on a common principle: human oversight should be proportionate to the consequence and materiality of the decision. For higher-impact uses, a bank that keeps a competent person able to review, intervene in or override an automated decision before it takes effect satisfies most markets' binding position. The two that restrict solely automated decisions add a step of their own, so meaningful oversight is common ground in most of the region.

Supervisors and banks describe this in the same terms. One supervisor put the principle plainly, that **"technology remains neutral and that ultimately the person is responsible."** A bank described keeping **"strict human-in-the-loop validation before a higher-risk decision takes effect, with continuous monitoring for drift"**. The common ground here is already close to what a shared expectation would ask for.

FAIRNESS

Fairness is the requirement to guard against AI that produces biased or discriminatory outcomes. A model learns from historical data and can carry the patterns in that data into its decisions, at the scale and speed of automation. In a 2024 study, researchers at Lehigh University gave several widely used AI language models thousands of mortgage applications that were identical apart from the applicant's race, and the models recommended more denials and higher interest rates for Black applicants, who would have needed around 120 additional points of credit score to be approved at the same rate as white applicants. Instructing the model to avoid bias removed the racial gap in approvals and narrowed the gap in interest rates, which showed the bias came from the model itself.[69] An AI system can produce discriminatory outcomes even when no one intends it to, so a model's decisions have to be tested for bias before they are relied on.

This is a pillar where the region shares both a common foundation and a common gap. The foundation is that every market already imposes a binding duty to treat customers fairly, through its consumer-protection, market-conduct and responsible-lending rules, which is the closest existing counterpart to algorithmic fairness. Laos requires financial services to be provided without discrimination on stated grounds and imposes a responsible-lending and affordability duty,[70] and Thailand's responsible-lending rules set a binding affordability discipline for credit.[71] Several markets reinforce this at the data layer: the credit-reporting rules in Cambodia, among others, bar the credit bureau from holding data on attributes such as race, colour and political belief, which limits the scope for a downstream model to discriminate by proxy.[72] Cambodia's 2026 technology and cyber-risk guidelines now also ask banks to keep their AI systems fair and free of biased data. This is one example of the AI-specific fairness expectations beginning to appear across the region.[73]

The common gap is that hardly any market has an operational way to test whether an AI model is fair. A binding non-discrimination principle tells a bank what outcome to avoid, but few rulebooks set out the fairness metrics, disparate-impact tests or subgroup checks that would show whether a model meets it. Two markets have begun to close that gap, in different ways. Singapore has a developed, operational fairness methodology: its voluntary Veritas toolkit takes a bank through a structured assessment that asks it to declare its fairness objectives, measure performance against them, justify any use of personal attributes and their proxies, and monitor the system once it is in use. It was first built and tested on credit-scoring and marketing use cases, and later extended to insurance underwriting.[74] Its open-source toolkit reached a second version in 2023, with a refined fairness assessment.[75] It is deliberately a process, setting no numeric threshold a bank must clear and leaving banks to set their own, on the view that there is no single answer to what counts as fair.[76] Malaysia has taken the binding route: its fair-treatment rules require a bank to exercise due care when it adopts AI or machine learning in credit assessment or underwriting, and to avoid unfair discrimination, excessive pricing and the exclusion of vulnerable consumers.[77]

62 Philippines, Data Privacy Act of 2012 (RA 10173) and Implementing Rules and Regulations, 2016
63 National Privacy Commission, Advisory No. 2024-04, Guidelines on Artificial Intelligence, 2024
64 Vietnam, Law on Artificial Intelligence No. 134/2025/QH15, 2025
65 Indonesia, Law No. 27 of 2022 on Personal Data Protection, 2022
66 Personal Data Protection Commissioner (Malaysia), Automated Decision-Making and Profiling Guideline, 2026
67 Monetary Authority of Singapore, Guidelines on Artificial Intelligence Risk Management, 2026
68 Bank of Thailand, Policy Guideline on Risk Management of the Use of AI Systems, 2025
69 Bowen, Price, Stein and Yang, Measuring and Mitigating Racial Disparities in Large Language Model Mortgage Underwriting, Lehigh University, 2024
70 Lao PDR, Decree on Financial Consumer Protection No. 225/GOV, 2020
71 Bank of Thailand, Notification No. 3/2568 on Responsible Lending, 2025
72 National Bank of Cambodia, Prakas on Credit Reporting, 2020
73 National Bank of Cambodia, Technology and Cyber Risk Management Guidelines, 2026
74 Monetary Authority of Singapore, MAS-led industry consortium publishes assessment methodologies for responsible use of AI by financial institutions (Veritas Phase 2), media release, 2022
75 Monetary Authority of Singapore, Toolkit for Responsible Use of AI in the Financial Sector (Veritas Toolkit version 2.0), media release, 2023
76 Veritas Consortium (MAS and industry), Veritas Document 1: FEAT Fairness Principles Assessment Methodology, 2020
77 Bank Negara Malaysia, Fair Treatment of Financial Consumers, policy document, 2024

© 2026 GFTN Limited, All Rights Reserved. Reproduction Prohibited.

41%

of banks apply bias or fairness testing to higher-risk AI.

Source: the GFTN and ABA survey of ASEAN bank leaders.

In some markets, the firmer fairness obligations come from outside the banking rulebook. In the Philippines, the privacy regulator's AI advisory, which sets out how the data-privacy law applies to AI, requires organisations to identify and monitor bias in their AI systems, considering systemic, human and statistical bias.[78] The financial-consumer-protection statute directs that supervisory rules be guided by internationally accepted standards, which gives the central bank a ready basis to adopt a common fairness reference.[79] Vietnam approaches fairness on two levels. Its 2025 Law on Artificial Intelligence sets a binding statutory principle that AI activity be fair, transparent and free from bias and discrimination.[80] Its consumer law reverses the burden of proof, so a bank defending a model-driven outcome must show it was not at fault. The same law gives specific protection to defined categories of vulnerable consumer.[81]

Because fairness is an agreed principle across the region and the missing piece is a shared, practical way to test for it, it is one of the areas where a common reference would add the most. Singapore's methodology offers a candidate process and Malaysia's rule a candidate standard, and the Philippine statute shows how a market could take up such a reference through its existing commitment to international practice.

DATA

Data covers how customer information is governed and protected, and in particular whether it can move across borders. Banks hold some of the most sensitive information people have, and when that information is exposed, the harm is large and lasting. A single breach at one credit bureau in 2017 compromised the personal records of around 147 million people.[82] Data is also one of the two pillars on which some markets require a local implementation, which places it at the centre of any question of interoperability.

Reading note

WHAT THE DATA AND THIRD-PARTY PILLARS EACH COVER

The data pillar is about whether customer data may cross the border: how it is protected, and whether it can be sent abroad. The third-party pillar is about where the systems and data centres that hold it must physically sit, and how far a bank depends on a few shared providers.

A single rulebook can do both at once. A rule that data must be stored in the country is a location rule, so it sits under third-party; a rule that gates or blocks sending data abroad is a movement rule, so it sits under data. Where one market does both, it appears under each pillar for the part it governs. Filing each rule under the part it actually controls keeps one requirement from being read, or counted, as two. Read together, the two pillars describe how much control a bank keeps over its data and systems. This includes who can access and audit them, and whether the bank can change provider.

The baseline, though, is shared. Every market protects customer information, through banking-secrecy rules long embedded in its banking law and, increasingly, a modern data-protection regime. Most now have a dedicated data-protection law in force or recently enacted, among them Singapore's data-protection framework,[83] Thailand's Personal Data Protection Act,[84] and Brunei's data-protection order, whose main provisions took effect in 2026.[85] Privacy and data governance is also one of the seven principles of the ASEAN Guide,[86] so the region already shares the language of protection, if not yet all of its mechanics.

There is genuine convergence, too, on the mechanics of moving data across borders. A growing number of markets recognise the same instruments for transferring data under safeguards, and the ASEAN Model Contractual Clauses give the region a common tool of its own.[87] Thailand names both the ASEAN clauses and their European equivalent in a binding rule;[88] Malaysia's cross-border transfer guideline recognises the same family of instruments;[89] and the Philippines recognises the model clauses through its privacy regulator.[90] Where markets accept these shared instruments, a bank can move data between them under a recognised safeguard rather than negotiating each route from the beginning.

The differences lie in how far that recognition extends. Where a market has not adopted the shared instruments, it falls into one of two groups. In the first, data can leave without the regulator's sign-off, but the route is not yet settled. Brunei's data-protection order is in force, while the requirements for transfers abroad have not yet been prescribed.[91] In Laos, Myanmar and Timor-Leste, cross-border movement rests mainly on banking secrecy.[92,93,94] A bank in these markets establishes its own legal basis, such as customer consent, for each flow. In the second group, the regulator must clear the transfer itself and can refuse or later halt it. In Vietnam a

78 National Privacy Commission, Advisory No. 2024-04, Guidelines on Artificial Intelligence, 2024
79 Philippines, Republic Act No. 11765, Financial Products and Services Consumer Protection Act, 2022
80 Vietnam, Law on Artificial Intelligence No. 134/2025/QH15, 2025
81 Vietnam, Law on Protection of Consumers' Rights No. 19/2023/QH15, 2023
82 US Federal Trade Commission, Equifax to Pay $575 Million as Part of Settlement over 2017 Data Breach, 2019
83 Singapore, Personal Data Protection Regulations, 2021
84 Thailand, Personal Data Protection Act B.E. 2562, 2019
85 DLA Piper, Data protection laws in Brunei, 2026 (secondary source)
86 ASEAN, ASEAN Guide on AI Governance and Ethics, 2024
87 ASEAN, ASEAN Model Contractual Clauses for Cross-Border Data Flows, 2021
88 Personal Data Protection Committee (Thailand), Notification on Cross-Border Personal Data Transfer (Section 29), 2023
89 Malaysia, Personal Data Protection Guidelines No. 3/2025: Cross Border Personal Data Transfer, 2025
90 National Privacy Commission, Advisory No. 2024-01 on Cross-Border Transfers, 2024
91 Brunei Darussalam, Personal Data Protection Order, 2025 (S 1/2025), 2025
92 Lao PDR, Law on Commercial Bank (Amended) No. 39/NA, 2023
93 Myanmar, Financial Institutions Law No. 20/2016, 2016
94 Timor-Leste, UNTAET Regulation No. 2000/8 on Bank Licensing and Supervision, 2000

© 2026 GFTN Limited, All Rights Reserved. Reproduction Prohibited.

transfer of personal data is made by filing a dossier with the authority, which keeps the power to direct or restrict it.[95] Moving data classified as core, which includes bank-account and payment data above set thresholds, needs the Ministry of Public Security to clear an impact assessment before it leaves the country.[96] In Cambodia a binding banking-secrecy duty governs customer data, and the central bank's 2026 guidelines require prior approval before it moves offshore.[97] Several of these positions are still in motion as newly enacted laws are implemented, so the map is expected to keep shifting. In these markets, the central set-up has to adapt to each market's own transfer rules. If more markets chose to recognise a shared transfer mechanism, such as the ASEAN Model Contractual Clauses, banks would have fewer separate routes to manage.

66%

of banks name data localisation and transfer rules as a cross-border friction for AI deployment.

Source: the GFTN and ABA survey of ASEAN bank leaders.

The supervisors and the banks describe these rules as conditions that shape how AI is built and where it runs, while still allowing it. A bank operating across the region made the same point, that these requirements **"do not prevent AI adoption, though they call for additional controls, architecture adjustments or local processing."**

CYBER SECURITY

Cyber security is the protection of AI systems and the data behind them from attack and compromise. AI widens the attack surface in two directions at once: it gives banks new defences, and it gives attackers new weapons. In a 2024 review of the financial sector, the United States Treasury found that generative-AI tools let attackers run more targeted phishing and business-email compromise, build more capable malware, and help less-skilled attackers mount simple but effective attacks.[98] The same tools help an attacker impersonate a bank's own staff and customers, get past its identity checks and authorise payments, and can be turned on the AI models themselves. A cyber regime built for conventional technology risk has to stretch to cover them.

The rules here were largely in place before AI, and they follow a common template across the region. Most markets already run a binding cyber and operational-resilience regime for banks, built to international practice and applied to any critical system regardless of the technology inside it, so an AI system inherits those obligations without a new rule. The shared architecture is consistent: identify critical systems, protect and monitor them, report incidents, test defences through penetration and red-team exercises, and maintain business continuity and recovery. Singapore holds critical systems to hard limits on downtime, recovery time and incident notification;[99] Malaysia requires a security operations centre, red-team testing and ransomware-resilient recovery, and a breach may lead to enforcement action;[100] and Indonesia and Thailand run comparably concrete, maturity-based regimes with mandatory penetration testing and incident-reporting duties.[101,102] Because these frameworks are close to one another and to international practice, a bank meets them with the security controls it already operates.

In a few markets the cyber rules are lighter, resting on general business-continuity obligations or on a general cyber-crime law rather than a detailed banking cyber-security rulebook. The rules are thinner in those markets, but they do not conflict with the fuller ones elsewhere, so a bank can apply its own group security controls there without meeting a contradictory local requirement. Some markets have also begun to address the threats specific to AI. Singapore has set out emerging cyber threats arising from both the malicious use of generative AI and vulnerabilities in GenAI deployments, including deepfakes, AI-enabled phishing and malware, data leakage, and model manipulation, and outlines practical measures for financial institutions to mitigate these risks.[103] The Philippines requires the defensive use of AI: its rules oblige banks to use machine learning, alongside rule-based and other methods, in their fraud-management systems, and to keep those systems refined as fraud tactics change. This reverses the usual pattern, in which the rules place limits on how a bank may use AI; here a bank is instead obliged to deploy it.[104]

The region's languages add a dimension of their own. One authority ran a red-team exercise built for the banking sector and found that prompts in its own language could trigger jailbreaking more readily than English ones. A bank made a related point, that the region's many languages leave general-purpose models **"unreliable without costly localized fine-tuning."** Testing for these failures, and building models that resist them, is harder where the tools and benchmarks are built for other languages.

In three markets, cyber security also intersects with national-security law. In Myanmar, Vietnam and Laos the banking cyber regime now runs on two layers: the central bank's own binding technology-risk rules, and above them a critical-information-infrastructure framework administered by a national-security authority. Myanmar's 2025 cybersecurity law brings finance within that regime, with an annual cybersecurity report to the ministry and a power to require certain information to be stored in the country.[105] Vietnam's 2025 cyber law reaches into the bank's own operations, requiring an information-system owner to connect its security-monitoring systems to a state cybersecurity centre and to supply user information on short compulsion timelines.[106]

95 Vietnam, Personal Data Protection Law No. 91/2025/QH15, 2025
96 Vietnam, Decree 165/2025/NĐ-CP detailing the Law on Data, Article 12, 2025; and Prime Minister, Decision 20/2025/QĐ-TTg on the lists of important and core data, 2025
97 Cambodia, Law on Banking and Financial Institutions 1999, Article 47; and NBC Technology and Cyber Risk Management Guidelines, 2026
98 US Department of the Treasury, Managing Artificial Intelligence-Specific Cybersecurity Risks in the Financial Services Sector, 2024
99 Monetary Authority of Singapore, Notice FSM-N05 on Technology Risk Management, 2024
100 Bank Negara Malaysia, Risk Management in Technology (RMiT), policy document, 2026
101 Otoritas Jasa Keuangan, Circular Letter No. 29/SEOJK.03/2022 on Cyber Security and Cyber Resilience for Commercial Banks, 2022
102 Bank of Thailand, Notification No. 5/2566 on Information Technology Risk Supervision, 2023
103 Monetary Authority of Singapore, Cyber Risks Associated with Generative Artificial Intelligence, information paper, 2024
104 Bangko Sentral ng Pilipinas, Circular No. 1213 implementing the Anti-Financial Account Scamming Act, 2025
105 Myanmar, Cybersecurity Law (Law No. 1/2025), 2025
106 Vietnam, Law on Cybersecurity No. 116/2025/QH15, 2025

© 2026 GFTN Limited, All Rights Reserved. Reproduction Prohibited.

Vietnam's 2026 implementing decree, now in force, requires domestic providers of telecommunications, internet and value-added telecommunications services to store specified user data in Vietnam. It defines those services by reference to telecommunications law, so a bank's own services appear to fall outside it, and the State Bank's rules continue to allow a bank's primary and backup systems to sit offshore.[107] Laos's Law on Cybersecurity, published in March 2026, names banking as critical infrastructure, requiring banks to connect to a national cyber-security system and to report incidents to the technology ministry.[108] These are national security obligations that a bank's ordinary group security function is not usually designed for, such as connecting monitoring systems to a state facility or meeting government demands for information within tight deadlines. A bank operating across the region has to meet them as an extra, market-specific requirement, in addition to the cyber controls it applies everywhere else.

For cross-border operation, then, cyber security raises few obstacles. The banking resilience baseline is already shared and close to international practice, the early work on AI-specific threats is converging on the same concerns, and the main market-specific requirements a regional bank has to accommodate come from national-security regimes. A common reference on cyber resilience would build on ground the region has largely already reached.

THIRD-PARTY AND CONCENTRATION RISK

Third-party and concentration risk is the management of a bank's reliance on a small number of shared cloud, model and technology providers. AI deepens that reliance, because the models, computing power and specialist tools behind it are supplied by a handful of firms. The danger is that a single provider becomes a single point of failure for many banks at once. In July 2024 a faulty software update from a single cyber-security provider disabled millions of Windows computers worldwide, disrupting banks' cash machines, mobile applications and call centres alongside airlines and hospitals.[109] The more banks depend on the same few providers, the more a fault in one of them can propagate across the system at the same moment.

Most markets manage this reliance through an outsourcing-risk framework, and these frameworks share a common architecture. Wherever a bank hands a function to an external provider, it stays accountable for it, conducts due diligence, secures audit and access rights, controls sub-contracting, and holds an exit plan with the right to recover or delete its data. Timor-Leste and Myanmar place board oversight and internal-audit duties over outsourced functions but impose no dedicated due-diligence, audit-rights or exit-plan regime,[110,111] and Laos requires the central bank's approval and local record-keeping without the fuller framework.[112]

Cloud services are generally treated as a form of outsourcing and fall under the same rules. A typical regime rests on a binding notice and supervised guidelines that give the supervisor audit rights over the provider and make the bank's accountability non-delegable.[113] The regimes in Malaysia, Thailand and the Philippines take a similar shape, which makes cross-border consistency easier.

The requirement that calls for a local implementation concerns where a bank's systems and data may physically sit. These location rules sit in the outsourcing regime, so the report treats them here. Some markets require a measure of onshore residency. Indonesia requires a bank to keep its electronic systems, data centres and recovery centres inside the country by default, and to obtain a regulatory licence, backed by a no-objection from the provider's home supervisor, before it places or processes them offshore.[114] The regulation names six purposes for which systems may be placed offshore on these terms. These include, for instance, analysis required by a regional banking group's home authority, group risk management, group anti-money-laundering work and services to global customers. The other two cover communication within the group and the bank's internal management. Indonesia's 2026 technical rule carries this requirement forward, and lists knowledge-management tools that use AI among the internal-management systems that may sit offshore.[115] Cambodia requires a bank to keep at least one primary data centre in the country to host its IT environment, and allows a data centre abroad only as a secondary site, with the central bank's approval. The central set-up therefore cannot act as the primary host for a bank in Cambodia, so part of it has to run in the country. The central bank's prior approval is also needed before an enabling technology such as cloud runs a critical system.[116] Thailand sets a lighter requirement of the same family: the backups of significant cloud systems must be held in Thailand or in a country other than the primary one, which still permits offshore storage.[117] Markets also divide over whether adopting cloud or a material outsourcing needs prior regulatory approval or only advance notification.

Several markets have shaped their rules to recognise the operating model that a central set-up may rely on. The Philippines maintains a formal category for banks whose critical processing sits offshore within their own group, permits public cloud even for core banking against an internationally recognised assurance report, and exempts intra-group arrangements from prior approval where the provider is a regulated group entity, including when the arrangement is offshore.[118] Malaysia exempts an intra-group outsourcing from prior approval where the affiliate is regulated and the two supervisors cooperate,[119] and Vietnam places intra-group technology provided to the local branch of a foreign bank outside its outsourcing regime altogether.[120] These carve-outs mean the group operating model is one that several rulebooks already accommodate. A supervisor from another market cautioned against **"the wrong impression"** that **"the**

[107] Vietnam, Decree 333/2026/ND-CP detailing the Law on Cybersecurity, Articles 3 and 19, 2026; and State Bank of Vietnam, Circular 09/2020/TT-NHNN on information system security, Articles 35 and 51
[108] Asia IP, Laos publishes new law on cybersecurity, 2026 (secondary source)
[109] IBM, Recent CrowdStrike outage: what you should know, 2024
[110] Banco Central de Timor-Leste, Instruction No. 29/2025 on the Corporate Governance of Banks, 2025
[111] Central Bank of Myanmar, Risk Management Practices of Banks Guideline, 2020
[112] Lao PDR, Law on Commercial Bank (Amended) No. 39/NA, 2023
[113] Monetary Authority of Singapore, Notice 658 on Management of Outsourced Relevant Services for Banks, 2024
[114] Otoritas Jasa Keuangan, Regulation No. 11/POJK.03/2022 on the Implementation of Information Technology by Commercial Banks, 2022
[115] Otoritas Jasa Keuangan, Board of Commissioners Regulation (PADK) No. 1 of 2026 on the Implementation of Information Technology by Commercial Banks, 2026
[116] National Bank of Cambodia, Technology and Cyber Risk Management Guidelines, 2026
[117] Bank of Thailand, Notification No. 5/2566 on Information Technology Risk Supervision, 2023
[118] Bangko Sentral ng Pilipinas, Manual of Regulations for Banks, Section 148 (IT Risk Management, consolidating Circular No. 1137), 2022
[119] Bank Negara Malaysia, Outsourcing, policy document (BNM/RH/PD 028-93), 2019
[120] State Bank of Vietnam, Circular 09/2020/TT-NHNN on information system security, 2020

© 2026 GFTN Limited, All Rights Reserved. Reproduction Prohibited.

development of AI in [its market] is restricted", explaining that whether a system may run offshore depends on what it is used for and whether the data is individual or aggregated.

Each market's rules look closely at the single arrangement in front of them: this bank, using this provider, for this service. They make sure that arrangement is sound and that the bank could exit it if it had to. What they do not capture is the wider pattern that forms when many banks make similar choices and come to depend on the same few cloud and AI providers at once. That build-up is concentration risk, and it is what would turn one provider's failure into a problem for a large part of the banking system. Because each bank's reliance is judged on its own, this system-wide pattern falls outside what the rules are built to assess.

Third-party governance in the region therefore has two sides. Deal by deal, the rules are broadly similar from one market to the next. Some markets have deliberately shaped them to accommodate banks that run shared technology across a regional group. Running a central set-up across the region is therefore workable, although data-residency rules shape where parts of it have to sit. Concentration is the harder side. The same handful of providers now sits beneath banks in several countries at the same time, so the exposure builds up across the region as a whole. A risk that spans borders in this way is a natural subject for supervisors to work on together, rather than one that each national rulebook can manage on its own.

The banks and their supervisors both see this building up. In the bankers' survey, 61% of banks rate dependency on a few external foundation-model or cloud providers as a high or medium operational-resilience risk, and only 10% see no risk in it at all. One supervisor described the exposure directly, warning that with banks leaning on **"two or three big providers, the same problem in one bank easily affects the others."** A standing regional view of where that dependence concentrates is what would bring it into sight.

61%

of banks rate dependency on a few external foundation-model or cloud providers as a high or medium operational-resilience risk.

Source: the GFTN and ABA survey of ASEAN bank leaders.

READING THE PILLARS TOGETHER

Read across all eight pillars, a single pattern stands out. On most of them the region already shares the ground a common approach would need. Where it does not, the difference is usually an overlay the central set-up can absorb, or a gap the whole region shares and could close together. Only a few requirements call for a local implementation. Some markets require a bank's primary systems, and the data they hold, to be hosted in the country. In others, transferring customer data abroad is subject to the regulator's approval. A bank has to plan for the case where a transfer is not approved, so it needs storage and processing for that data in the country. Even in these cases, a bank can keep its regional approach. It establishes and runs the affected part in the country, and keeps it aligned with its central set-up.

One difference runs beneath all eight pillars. The markets do not yet share a single definition of artificial intelligence, a common way to sort its uses by how much risk they carry, or a settled view of which institutions and which techniques their rules should reach. A binding legal definition is rare, and more often the term is used without being pinned down. This shared vocabulary is the groundwork any wider alignment would rest on, because two markets can only recognise each other's rules once they mean the same thing by the words in them. Exhibit 6 sets out, pillar by pillar, how the eleven markets relate to a shared regional approach.

© 2026 GFTN Limited, All Rights Reserved. Reproduction Prohibited.

Exhibit 6:

How the eleven markets relate to a shared regional approach, pillar by pillar

Pillar	Common ground	Overlay	Local implementation	What this means
Governance and accountability Clear responsibility for how AI is used sits with the board and senior managers.	11			Every market already makes the board and senior managers legally responsible for what the bank does, so AI decisions are covered by rules that already exist. Some markets go further and assign accountability to named senior officers, which AI attaches to.
Model risk Testing and validating an AI model before relying on it, and keeping it under review in use.	9	2		Most markets test AI models under their general risk rules, and some add detailed guidance. There is no single standard the region shares for validating a model before a bank relies on it. **Overlay (two markets):** One makes independent model validation and annual model-risk reporting binding; the other requires an AI system on a critical function to prove a low error rate before it runs with less human involvement.
Fairness Guarding against AI that produces biased or discriminatory outcomes.	11			Every market requires customers to be treated fairly, the shared foundation. What the region does not yet share is a common way to test whether an AI model is fair; where markets have gone further they have done so in two ways, a structured testing process or a binding fair-treatment duty for AI in credit.
Cyber security Protecting AI systems and the data behind them from attack.	8	3		Broadly similar cyber-security rules are in place across the region, though with varying reach, and an AI system is covered by them. Work on the threats specific to AI is converging on the same concerns. **Overlay (three markets):** Three markets run the bank cyber regime on two layers, the central bank's own binding technology rules and, above them, a national-security framework. One links security monitoring to a state centre and compels fast handover of user data; another requires annual reporting to the ministry and can require some data kept in-country; the third requires banks to connect to a national cyber-security system and report incidents to the technology ministry.
Human oversight Ensuring a human is able to review, step into or overturn an AI-driven decision.	9	2		Nearly every market gives customers a binding route to challenge a decision and have it reviewed, through complaints and redress, and on AI the markets converge on keeping a competent person able to review and override an automated decision before it takes effect. Few markets yet give a customer a specific right to contest an automated decision after it is made. **Overlay (two markets):** one bars a solely automated decision with legal effect unless the customer consents; the other gives a right to object to a decision based solely on automated processing that has a significant effect.

© 2026 GFTN Limited, All Rights Reserved. Reproduction Prohibited.

Exhibit 6:

How the eleven markets relate to a shared regional approach, pillar by pillar (continued)

Pillar	Common ground	Overlay	Local implementation	What this means
Transparency Being able to explain what an AI system does, and to disclose when it is being used, so a decision can be questioned.	4	7		Every market requires products to be clearly disclosed, but the duty to explain an AI decision is newer, mostly non-binding, and in several markets set by privacy regulators. **Overlay (seven markets):** Markets have varying requirements, including explaining what a model weighs, disclosing when a decision is automated or when a customer is dealing with a chatbot, and explainability expectations set in AI guidance.
Third-party and concentration risk Managing a bank's reliance on the few cloud, model and technology providers behind AI, so one provider's failure cannot ripple across many banks.	8	1	2	Rules on using outside technology providers are broadly similar, and cloud is treated the same way. Banks lean on the same few foundation-model and cloud providers, a system-wide exposure no market yet manages, only bank by bank. **Overlay (one market):** Requires backups of significant cloud systems to sit in-country or in a third country, which still allows offshore storage. **Local implementation (two markets):** Both require a bank's primary systems to sit in-country: one unless the regulator licenses offshore placement for defined purposes, the other allowing offshore sites only as secondary. The bank runs this part of its set-up in each of these markets rather than centrally.
Data How customer data is governed and protected, and in particular whether it can move across borders.	5	4	2	Every market protects customer data, through banking secrecy and, in most, a data-protection law. Five also recognise standard mechanisms for moving it abroad. **Overlay (four markets):** Data can leave without the regulator's sign-off, but the route is not yet settled. In some markets the transfer rules are not yet in operation. In others, banking secrecy leaves unclear whether a group company or offshore provider may receive customer data. The bank establishes its own legal basis, such as customer consent, for each flow. **Local implementation (two markets):** Transferring customer data abroad is subject to the regulator's approval, so banks plan to store and process it in the country. One requires security-ministry clearance before large volumes of customer data leave the country. The other requires central-bank approval before customer data is hosted abroad.

Numbers count the eleven markets in each band, as at October 2026.

- **Common ground:** the central set-up operates under the rules as they stand
- **Overlay:** met within the central set-up through an additional step
- **Local implementation:** part of the set-up is established and run in the market

Source: GFTN analysis of the eleven country profiles.

© 2026 GFTN Limited, All Rights Reserved. Reproduction Prohibited.

06

CASE STUDY

A regional platform in practice

Note

The bank and the markets in this case are hypothetical. The purpose is to show how the pattern set out in the previous section, common ground, overlays and local implementations, plays out for a group that builds one AI system and extends it across the region. It illustrates where a shared approach would help; it does not prescribe one.

A regional banking group sets out to tackle a familiar problem. Its financial-crime teams are seeing more fraud that moves across borders: an account opened in one market receives a payment from a second, sends funds through a third, and connects to accounts that have already raised alerts elsewhere. The group builds an AI-enabled monitoring platform to bring these signals together, so that investigators can see unusual transaction chains, detect likely mule accounts and prioritise the highest-risk cases. Its governance design is clear: a central team validates the model, local teams oversee deployment, and specialists review material alerts before any customer is affected. The group then sets out to extend the platform across all of its ASEAN operations.

Across most of what the platform must satisfy, the group is on common ground. The expectations that the board is accountable for what the system does, that a person reviews a material alert, that the model is tested before it is relied on, and that the data behind it is protected, are broadly shared from one market to the next. On these the group builds once, to a single standard, and the design travels. The common ground, though, stops short of automatic interoperability. Even where two markets require the same independent model validation, each supervisor still expects its own evidence, so the group runs a separate validation exercise in each place for a control everyone agrees on. Mutual recognition, a thorough check in one market counting for more in the next, is what would let one exercise carry further.

In several markets an additional requirement sits on top of the shared expectation. One market requires the model to be validated and reported each year and cleared before it continues in use. Another gives a customer a right to object to a fully automated decision. A third sets specific disclosure expectations. None of these calls for a separate platform, and the group absorbs each as a step in the regional operating model. What it cannot avoid is that the same model acquires a local governance file in each market, because each supervisor classifies and documents it in its own way. A shared way to classify AI use cases is what would let one account of the model do more of the work.

The platform's value is the regional view. A fraud pattern that surfaces in one market often runs through accounts and counterparties in the next, and the group wants to follow it across the markets it touches.

The harder constraint is data. For sharing specific financial-crime intelligence across borders, the group can often use established channels built for that purpose, and within any single market it can share more directly still with other banks and its supervisor. These channels, though, do not let it pool all of the customer and transaction data the models run on into one place. In some markets the models run on local data, and the cross-border picture is assembled by coordination between the group's teams working from their own local views. This is where a local implementation goes beyond an overlay. Here, part of the system has to run separately in those markets.

A last issue appears only at scale. As the platform spreads, the group can see that several of its markets now depend on the same small set of cloud and technology providers. Each local entity has done its own due diligence, and each arrangement meets the applicable rules, yet no single review shows how much of the group's regional activity rests on the same few suppliers, or what a prolonged outage at one of them would do across markets at once. A standing regional view of where this dependence concentrates is what would bring it into sight.

By the end of the rollout, the group has one platform, but it runs differently from one market to the next. Some markets run it in full. Others run a local variant, with extra manual steps and more cautious limits. The aim was the same everywhere: AI that is accountable, reliable and well supervised. The differences lay elsewhere: the extra requirements a few markets added, the evidence each supervisor wanted for the same controls, and the limits on moving data across borders. The next section outlines opportunities for interoperability that would help a regional bank like this one build more of its platform once, so it repeats less of the work in each market.

© 2026 GFTN Limited, All Rights Reserved. Reproduction Prohibited.

07

INTEROPERABILITY OPPORTUNITIES

Common standards, shared tools and supervisory alignment

Section 5 showed where the region's rules converge and where they differ. This section sets out ten emerging opportunities for consideration. They are illustrative. They are shared to show what is possible, and selecting priorities and defining outcomes would be part of the next phase of work. The opportunities fall into three types. Common standards are shared reference points that banks and supervisors could agree to use. Shared tools and skills are practical resources that banks and supervisors could use together. Supervisory alignment covers steps that need deeper agreement and trust between supervisors. The first two build on ground the region already shares, and they could start in the near term. Supervisory alignment asks more, and it would follow.

The three descriptions of fit set out in Section 3 each suggest a different response. Where markets share common ground, banks and supervisors could standardise where appropriate, or jointly develop components that every market can reuse. Where a market adds an overlay, the region could agree a common minimum baseline, with each market's own requirements sitting on top of it. Where a requirement calls for a local implementation, the aim would be to accommodate and manage the difference. These requirements reflect each market's own reasons of security and sovereignty, so harmonising them may be neither practical nor appropriate. These principles are indicative, and how they apply to each opportunity would be settled in the next phase of work.

Several of the opportunities concern customer data and customer rights. In most markets these fall to the data-protection authority as well as the banking supervisor, so both would need to take part. For each opportunity, the report sets out why it matters, what it could involve, and where a comparable approach already works. The examples show what an opportunity could look like in practice.

As Exhibit 7 shows, the opportunities would help banks in different ways: for banks operating across ASEAN, the main benefit is efficiency, and for banks focused on a single market, a clear starting point. Supervisors would also gain, as they could draw on shared methods and evidence and see risks that cross borders earlier. Annex B sets out, for each opportunity, the evidence behind it and the banks it would help most.

This matters because the markets are at different stages. Some have set detailed expectations for AI, and others have yet to write rules for AI in banking. A common baseline gives the second group a reference to build from. It gives the first group a way to make its rules easier to read across borders.

COMMON STANDARDS

A shared vocabulary and governance baseline. The markets do not share common definitions of artificial intelligence or a common way of describing its uses. Few markets define AI in binding law, and many rules use the term without defining it. Many banks are also still putting their own AI governance in place. In the bankers' survey, 49% of banks reported having no internal AI governance policy in force in any of the seven areas of the ASEAN Guide on AI Governance and Ethics. The gap is wider for banks that operate only in their home market. On average, banks without cross-border operations have a policy in force in 25% of the seven areas, against 42% for banks with them. A shared vocabulary and a governance baseline are a practical place to start, since rules are easier to read across borders when the terms behind them mean the same thing from one market to the next.

This could involve agreeing a common set of terms and a baseline for the governance expectations all eleven markets already share. Both would build on the ASEAN Guide,[121] which gives the region a foundation, even if its terms are not yet used consistently across the markets' rules. They could also draw on the AI guidance that individual markets have issued. Each authority would align the terms in its own rules with the agreed definitions. Banks could use the baseline as a reference for their own policies, roles and documentation. The aim is a common reference each market can adopt without rewriting its own rulebook.

A shared definition already operates internationally. OECD member countries approved a common definition of an AI system in 2023,[122] and the European Union aligned the definition in its AI Act to it,[123] so different regimes describe the same thing in the same way. The region would apply the same approach, anchored in its own Guide.

The region's own supervisors have asked for this. One supervisor saw value in **"markets aligning on AI taxonomy, AI risk management, and how to implement it in each use case"**. Another described how a baseline could work: **"A principles-based approach, rather than very prescriptive rules, would be most helpful, particularly on the governance and accountability of AI. A minimum baseline could be pitched as a must-have, which every country has anyway, with certain parts set as best practices that countries, and banks, could strive towards."** A bank operating across several markets said that differing supervisory definitions of AI oblige it to keep its own taxonomy simply to reconcile them. In a roundtable discussion with industry participants, one said **"we all are invested in making sure we're using the same taxonomy and language."**

121 ASEAN, ASEAN Guide on AI Governance and Ethics, 2024
122 Organisation for Economic Co-operation and Development, Updated OECD definition of an AI system, 2023
123 European Union, Regulation (EU) 2024/1689 (Artificial Intelligence Act), Article 3 and Recital 12, 2024

© 2026 GFTN Limited, All Rights Reserved. Reproduction Prohibited.

Exhibit 7:

Emerging interoperability opportunities: common standards, shared tools and supervisory alignment

NEAR-TERM		LONGER-TERM
Common standards Shared reference points that banks and supervisors could agree to use	**Shared tools and skills** Practical resources that banks and supervisors could use together	**Supervisory alignment** Steps that need deeper agreement and trust between supervisors
1 **A shared vocabulary and governance baseline**	4 **Shared tools for testing AI systems**	7 **A shared risk tiering for AI use cases**
2 **A shared baseline for transparency to customers**	5 **A practical playbook for moving data across borders**	8 **Portability of model validation**
3 **A shared disclosure standard for foundation-model providers**	6 **Capacity building**	9 **Cross-border sharing of AI incidents**
		10 **Regional visibility of provider concentration**

Regional Across ASEAN	For banks operating across ASEAN, the principal benefit is efficiency. Common definitions, testing standards and evidence could be reused across markets. Closer supervisory alignment would reduce duplicate validation and approvals, while clear cross-border data pathways would address their most persistent operational constraint.
Local Within each market	For banks focused on a single market, the principal benefit is a clear starting point. A common baseline offers a reference to adopt while national rules mature. Shared tools and training provide access to expertise that would be costly to develop independently.

Source: GFTN analysis.

A shared baseline for transparency to customers. AI increasingly informs decisions about customers, such as whether a loan is approved. Transparency is the pillar on which the markets add the most requirements of their own: seven of the eleven markets add an overlay (Exhibit 6). The banks want more clarity here. In the bankers' survey, 69% of banks named transparency and explainability among the three areas that would benefit most from greater regulatory clarity, second only to privacy and data governance. Just over a third of banks (34%) have an internal transparency and explainability policy in force. One supervisor saw value in **"guidance or practical examples of things like explainability and transparency"**, noting that the level of detail depends on the audience: **"Transparency to whom? To customers or to supervisors, and how technical or how plain you should be."**

This could involve agreeing a common baseline in three areas. The first is what a bank tells customers when AI is used in a decision about them. The second is how it explains an adverse decision. The third is how a customer can challenge a decision that AI has driven. The baseline would need to leave room for local norms. Markets differ in how much customers expect to be told about AI, and in how they view AI decisions and the safety of their data. Each market's requirements could then be mapped against the baseline, so a bank can see where a local step is still needed. One customer-facing design could then work across markets, with local additions where a market requires them.

The European Union applies a common baseline of this kind across its member states. Its AI Act requires that people be told when they are interacting with an AI system. It also gives a person affected by a decision based on a high-risk AI system, such as one used to assess creditworthiness, the right to a clear and meaningful explanation of the system's role in the decision.[124] The same rules apply in every member state.

[124] European Union, Regulation (EU) 2024/1689 (Artificial Intelligence Act), Articles 50 and 86 and Annex III, 2024

© 2026 GFTN Limited, All Rights Reserved. Reproduction Prohibited.

In a roundtable discussion with industry participants, one noted that when an application is declined, the customer has a right to know why.

A shared disclosure standard for foundation-model providers. In ASEAN, foundation models are already in widespread use across banks: around three-quarters have generative AI in production or piloting, and the models behind them come from a small number of external providers. Each bank and each supervisor assesses those providers separately, so the same provider is reviewed many times over, each time against a slightly different set of requirements. In the bankers' survey, 39% of banks selected model or vendor approval as a factor that causes the most cross-border friction for AI deployment. One bank with branches in several markets said that consistency **"would reduce duplicated model validation and vendor-approval work across our overseas branches."** A disclosure can also go out of date. One authority noted that a provider's data-retention terms can change without notice.

This could involve the markets agreeing a common baseline for what a provider discloses to any bank or supervisor that assesses it. What the baseline covers would be for the parties to settle. It might include the model's capabilities and limitations, its data practices, and its governance and security. A provider could complete the disclosure once and keep it current, giving notice of material changes. A common baseline would reduce duplicated reviews. Local regulators and banks would still assess whether a model's performance, language coverage, data practices and limitations suit the local context.

Pharmaceuticals offer a parallel. Under the Drug Master File system, a supplier files its confidential information once with the regulator, and many manufacturers reference that single file without the supplier disclosing it to each of them.[125] One disclosure, held to a common standard, serves many assessors at once.

SHARED TOOLS AND SKILLS

Shared tools for testing AI systems. Common methods and tools for testing AI systems would let banks and supervisors compare results across markets. Testing an AI model for problems such as bias, drift and instability calls for specialist methods, and the ability to apply them is uneven across the region. In the bankers' survey, banks in the ASEAN-6, the six states that were members of ASEAN before 1995, reported applying bias or fairness testing to higher-risk AI far more often than banks in other markets (67% against 16%, Annex A). Language adds a further dimension. In discussions with regulators, one authority ran a red-teaming exercise with its banks because global models had not been tested in its language. It found that local-language prompts triggered jailbreaks more easily than English ones, and it shared example prompts with the banks. Results are also hard to compare when no two parties test in quite the same way. Shared tools would let any bank or supervisor test to a common standard.

This could involve pooling the assessment methods and open toolkits that already exist in the region, so they are available to any market or bank that wants them. It could include reference test scenarios and shared test sets for common use cases, including in the region's languages. It could also set out a common approach to when testing happens. Before a system goes live, tests could cover its accuracy, explainability and bias, and the human controls around it. Once it is in use, monitoring could cover drift and audits of human review. Banks could also test that they can switch an AI component off quickly, keep the underlying service running safely, and assess what the AI affected before it is restored. A common set of tests would also make results readable across markets, because two supervisors running the same test are looking at the same thing.

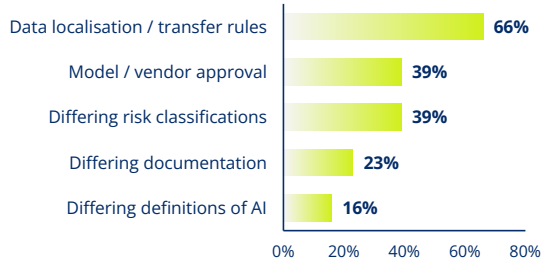
This kind of pooling already works internationally. The OECD maintains a public Catalogue of Tools and Metrics for Trustworthy AI, an openly available collection of methods for assessing AI systems that bodies and practitioners from many jurisdictions contribute to and draw on rather than each building their own. The region could assemble a common set in the same way.[126] One supervisor interviewed for this report named sharing methodologies for technical testing as an opportunity for regional alignment.

A practical playbook for moving data across borders. Data localisation and transfer rules are the largest source of cross-border friction for AI deployment that banks report. In the bankers' survey, 66% of banks selected them as a factor that causes the most cross-border friction, well ahead of any other factor. This includes 40% of banks without cross-border operations. One bank noted that the models it can use are **"effectively limited to those available domestically."** Another said that **"data localization is shaping infrastructure, not just compliance."** Three supervisors raised AI in fraud and anti-money-laundering work without being asked, and that work crosses borders by nature. In several markets a bank cannot draw on a standard, pre-approved transfer route. Some gate cross-border transfers case by case, and others have no settled data-transfer framework yet.

Exhibit 8:

Factors that cause the most cross-border friction for AI deployment

Share of the 61 surveyed banks that selected each factor; banks could select more than one.



Source: the GFTN and ABA survey of ASEAN bank leaders.

This could involve a playbook that sets out, market by market, the lawful routes for transferring data and the local requirements that apply. Because markets take different approaches, the playbook could offer options for different types of data, and each market could choose the ones that suit it. For some data, only results, summaries or anonymised information would cross borders, while the data itself stays in the country. Other data could move with safeguards where national law allows. To support these routes, the

125 US Food and Drug Administration, Drug Master Files (DMFs), 2025
126 Organisation for Economic Co-operation and Development, Catalogue of Tools and Metrics for Trustworthy AI, accessed September 2026

© 2026 GFTN Limited, All Rights Reserved. Reproduction Prohibited.

playbook could include standard templates for assessing a transfer and its recipient, which banks could reuse within existing law. It would also show where existing contractual mechanisms, such as the ASEAN Model Contractual Clauses,[127] can already be used. Markets that want easier movement could build on these by recognising the clauses between themselves. The playbook could also draw on the ASEAN Banking Interoperable Data Framework.[128] Throughout, the playbook would work within each market's rules. Where a market requires data to stay within its borders, for reasons of security and sovereignty that are its own to weigh, the playbook would not change that. The Digital Economy Framework Agreement could give this work a regional anchor (see Section 8).

The mechanism is proven. The European Union's Standard Contractual Clauses are a standard, pre-approved set of terms that give firms a lawful basis to move personal data across borders without renegotiating each transfer, and they are used at scale.[129] The region has its own version in the ASEAN Model Contractual Clauses, modelled on that approach and mapped to the EU clauses through a joint ASEAN and EU guide. Their reach depends on how widely the markets recognise them as a basis for transfer.

Capacity building. Governing AI well takes people, expertise and infrastructure, and these are not evenly distributed across the region. When each market builds its capacity alone, progress is uneven, and smaller markets bear costs that could be shared. Banks describe the gap in their own terms. One noted that its market's **"AI, data-governance and model-risk capability is still developing"**, and that **"regional guidance, shared standards and practical capacity-building"** would help it align more quickly. Another pointed to the **"limited availability of specialised AI governance talent and model-validation experts"** in its local market. Others point to a shortage of large, locally representative datasets, which limits how far models can be developed and validated at home. Several supervisors said they are still building their own AI expertise.

This could involve three kinds of support. Shared training could serve supervisors and bank staff. Regional expertise in model validation and AI governance could be made available to markets that lack it. Shared infrastructure, such as synthetic datasets for testing and training, could be built once and used across the region, subject to privacy rules. A market with fewer resources could then draw on the region rather than start from nothing. Annex A (Exhibit 9) shows that pooled capacity would add most where banks are still building formal model controls, such as bias testing and red-teaming. On the data itself, a bank that can draw on data from across the markets it serves is no longer confined to the data available in one market. Where real data cannot leave the country, synthetic data generated locally can serve the same purpose.

Central banks already pool innovation capacity this way. The Bank for International Settlements runs an Innovation Hub, set up in 2019, that develops shared technology as public goods and connects a network of central-bank experts, so members draw on common work rather than each building alone.[130] The same model would apply to AI-governance capacity in the region.

SUPERVISORY ALIGNMENT

These opportunities would require more deliberate coordination between markets and more trust between supervisors. Some also touch on national sovereignty. They are longer-term steps.

A shared risk tiering for AI use cases. The markets sort AI use cases by risk in different ways, and some are starting to set their own risk tiers. One rates banks' AI models as high, medium or low risk. Another's draft AI law proposes four tiers. In the bankers' survey, 39% of banks selected differing risk classifications as a factor that causes the most cross-border friction for AI deployment (Exhibit 8). When the same use case is treated as high-risk in one market and routine in another, a group operating across both must either meet the stricter requirement in both or configure the system differently in each. In a roundtable discussion with industry participants, some said most of their AI uses are simple and could be over-regulated without an agreed view of what counts as high risk. Common risk tiers would let one system meet a single standard across markets.

This could involve agreeing a shared set of risk tiers for AI use cases. Each tier would set out the minimum controls and testing expected, with illustrative classifications for common banking use cases. The work could also consider how generative and agentic AI fit within the tiers. It is a longer-term step, because it would ask markets to agree how much scrutiny a given use case warrants, a judgement that rests with each supervisor. The tiers would be a shared reference, and each supervisor's own classification would stand.

Common reference points for classifying AI already exist. The OECD Framework for the Classification of AI Systems, published in 2022, gives a shared set of dimensions for describing an AI system.[131] It is designed to help policy makers, regulators and legislators characterise AI systems in the contexts where they are used, since different systems bring different benefits and risks.

Portability of model validation. A banking group that validates an AI model in one market generally cannot rely on that work elsewhere, because each supervisor assesses the model against its own requirements. The desk research found no mechanism for one supervisor to rely on another's assessment of an AI model. The same model is validated again for each market the group operates in. In the bankers' survey, after data rules, the factors banks most often selected as causing the most cross-border friction were model or vendor approval (39%), differing risk classifications (39%) and differing documentation (23%). On average, banks rate the cost of meeting AI-related compliance across the markets they operate in at 3.5 out of 5.

This could start with agreeing what a sound validation covers, so that markets are describing the same exercise. A common evidence package and a standard evaluation scorecard would make results comparable and reusable. Supervisors could then choose to accept reusable testing components, and later the validation results themselves for selected use cases. Each supervisor could use this evidence alongside its own testing. As supervisors gain confidence in one another's assessments, mutual

127 ASEAN, ASEAN Model Contractual Clauses for Cross-Border Data Flows, 2021
128 ASEAN Bankers Association, ASEAN Banking Interoperable Data Framework (IDF): Guidance Document, Version 1.0, 2022
129 European Commission, Implementing Decision (EU) 2021/914 on standard contractual clauses for the transfer of personal data to third countries, 2021
130 Bank for International Settlements, BIS Innovation Hub, 2019
131 Organisation for Economic Co-operation and Development, OECD Framework for the Classification of AI Systems, 2022

© 2026 GFTN Limited, All Rights Reserved. Reproduction Prohibited.

recognition could follow, step by step and on equal terms. Building that confidence takes time, so it is a longer-term step.

Banking supervisors already cooperate in this way for international banking groups. For each large cross-border banking group, the supervisors of the markets where it operates meet as a supervisory college to share information and coordinate their work on the group.[132] Each supervisor keeps its own powers. Validation portability would apply the same idea to AI models. As one bank put it, a common control set **"answers multiple regulators from the same evidence."**

Cross-border sharing of AI incidents. When an AI system fails or is attacked in one market, the lesson is often relevant to banks elsewhere that run similar systems. More than half of banks (57%) named security and safety among the areas that would benefit most from greater regulatory clarity. Today every market's incident reporting runs to its own national authority, on its own timetable. Deadlines range from one hour in one market to no fixed deadline in another. One authority noted that a single incident can involve two or three jurisdictions. In a roundtable discussion with industry participants, one said sharing what has gone wrong would help more than new rules.

This could start with a common taxonomy for AI incidents, so that markets describe the same event in the same way. A voluntary protocol could then let serious AI incidents, cyber-threat intelligence and lessons learned move between supervisors and banks, with the confidentiality that sharing of this kind needs. Over time, the shared record could support a periodic regional outlook on AI risk.

Financial-sector information sharing of this kind already runs across borders. The Financial Services Information Sharing and Analysis Center (FS-ISAC) is a member-owned community through which banks and other financial firms share cyber-threat intelligence internationally.[133] It covers cyber threats in general. AI-incident sharing would extend the same mechanism, a trusted channel for timely sharing.

Regional visibility of provider concentration. A small number of foundation-model and cloud providers now sit beneath banks across the region at the same time. In the bankers' survey, 61% of banks rate dependency on a few external foundation-model or cloud providers as a high or medium operational-resilience risk, and only 10% see no risk in it at all. Supervisors review each bank's reliance on a provider one arrangement at a time. A review of one arrangement does not show how many banks rely on the same provider. Many banks across the region now rely on the same few providers, and no single national supervisor can see the full extent of this.

This could involve a standing regional view of where banks cluster on the same providers, built from data that supervisors already collect, using a common data template. Concentration would then be visible early, before a failure exposes it. Any such view would need clear rules on who collects the data, who can see it, and whether it is shared only in aggregate.

Supervisors elsewhere have begun to address this collectively. Under the European Union's Digital Operational Resilience Act, the authorities designate as 'critical' the information and communication technology (ICT) providers that many financial firms depend on. These providers are placed under a shared oversight framework. The providers that a large part of the system relies on are then identified and supervised jointly.[134] A regional view of where dependence concentrates would be a lighter first step towards the same end.

None of these steps would start from a blank page. Across the region, markets are already reaching for the same international standards and the same regional reference points, often without coordinating with one another. Much of the direction is already set. The opportunities above are illustrative. They show where coordinated effort could build on that direction, and choosing among them is the work of the next phase.

One regulator put the case for a shared approach simply:

> "What we need is collaboration, especially among ASEAN countries. Do not reinvent the wheel."
>
> **A regulator, on the case for a shared regional approach.**

132 Basel Committee on Banking Supervision, Principles for effective supervisory colleges, 2014
133 Financial Services Information Sharing and Analysis Center, About FS-ISAC, accessed September 2026
134 European Supervisory Authorities, Designation of critical ICT third-party providers under the Digital Operational Resilience Act, 2025

© 2026 GFTN Limited, All Rights Reserved. Reproduction Prohibited.

08

THE CONVERGENCE ALREADY UNDER WAY

Much of the groundwork for the opportunities highlighted in Section 7 is already present. Working largely independently, the region's markets have adopted the same international standards and built or joined the same regional instruments. Rules drafted separately have, as a result, converged on a shared set of reference points.

COMMON INTERNATIONAL STANDARDS

The region's markets have drawn on a common set of international standards. For the governance of AI itself, ISO/IEC 42001, a standard from the International Organization for Standardization (ISO) and the International Electrotechnical Commission (IEC), defines how an organisation should manage an AI system across its life cycle, and it has become a widely used benchmark.[135] The OECD AI Principles[136] and the AI risk-management framework issued by the United States' national standards body[137] serve as further common references for what responsible AI requires. For the financial sector specifically, the FSB's work on AI gives supervisors a shared basis for assessing the risks.[138] None of these instruments was written for ASEAN. Because markets have drawn on them independently, however, their national expectations have come to align.

The supervisors describe this convergence in their own words. One noted that **"several international frameworks were directly referenced"** in developing its guidance, **"including the OECD AI Principles and ISO"**, and that **"across jurisdictions, the various frameworks are generally grounded in a broadly consistent set of underlying principles."** Banks reach for the same anchors when their own rules run out, working from the ASEAN Guide and international frameworks such as the EU AI Act.

REGIONAL INSTRUMENTS

The region has also developed instruments of its own. The ASEAN Guide on AI Governance and Ethics states principles the markets hold in common and provides a shared starting point written for their context. On data, the ASEAN Model Contractual Clauses give markets a common contractual basis for transfers across borders, and a subsequent joint guide has aligned those clauses with their European equivalents, extending their reach beyond the region.[139] More broadly, the markets have concluded negotiations on a Digital Economy Framework Agreement (see box). Supervisory practice moves in parallel: approaches first developed by one market's authority have been adopted in the expectations of others, so alignment also spreads by example.

One supervisor summed up the state of play, seeing **"a significant degree of convergence in the underlying principles and no reluctance on our side to cooperate across the region. The guidelines converge; the detail is where the work remains."**

Regional development

THE ASEAN DIGITAL ECONOMY FRAMEWORK AGREEMENT

ASEAN Member States concluded negotiations on the Digital Economy Framework Agreement (DEFA) in May 2026.[140] Signing is planned for the 49th ASEAN Summit in November 2026, a commitment the region's economic ministers reaffirmed in September.[141] DEFA will be the region's first agreement dedicated to the digital economy.

Its scope extends to several areas relevant to banks. It covers digital trade, electronic payments and digital identity. It also covers data governance and privacy, and cooperation on cybersecurity. The framework that launched the negotiations set two objectives of particular relevance to this report. The first is to facilitate cross-border data flows while building frameworks for privacy protection. The second is to establish mechanisms for regulatory cooperation on emerging technologies, with artificial intelligence named as an example.[142]

For banking, DEFA could provide two tailwinds. It would give a regional anchor to the work on cross-border data transfers set out in Section 7. It would also establish a channel through which the markets could work towards a shared approach to AI.

Its effect on banking will depend on the final text, which had not been published at the time of writing. The region's earlier e-commerce agreement excluded financial services from its commitments on cross-border information flows and on the location of computing facilities.[143] Whether DEFA follows the same approach will become clear once its text is released. The agreement will also need to be ratified by Member States before it takes effect.

No regional body directed this convergence. Each market made its own choices, and those choices drew on the same references.

[135] International Organization for Standardization and International Electrotechnical Commission, ISO/IEC 42001:2023, Artificial intelligence management system, 2023
[136] Organisation for Economic Co-operation and Development, Recommendation of the Council on Artificial Intelligence (OECD AI Principles), 2019, updated 2024
[137] US National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework (AI RMF 1.0), 2023
[138] Financial Stability Board, The Financial Stability Implications of Artificial Intelligence, 2024
[139] ASEAN and European Commission, Joint Guide to ASEAN Model Contractual Clauses and EU Standard Contractual Clauses, 2024
[140] ASEAN, Statement of the Chairperson of the ASEAN Senior Economic Officials on the Conclusion of the ASEAN DEFA Negotiations, 2026
[141] ASEAN, Joint Media Statement of the 58th ASEAN Economic Ministers' Meeting, 2026
[142] ASEAN, Framework for Negotiating the ASEAN Digital Economy Framework Agreement, 2023
[143] ASEAN, ASEAN Agreement on Electronic Commerce, Article 7(4)(c) and 7(6)(c), 2019

© 2026 GFTN Limited, All Rights Reserved. Reproduction Prohibited.

THE COMMON GROUND ALREADY IN PLACE

The report's argument can now be drawn together. Section 5 established that the region's differences are narrower than they first appear. Across the eight pillars, most carry no conflicting rules. Where they differ, the differences are usually overlays that a shared approach can accommodate, or gaps common to the whole region that a single regional effort could close. Only a few requirements call for a local implementation. Some markets require a bank's primary systems, and the data they hold, to be hosted in the country. In others, transferring customer data abroad is subject to the regulator's approval. A bank has to plan for the case where a transfer is not approved, so it needs storage and processing for that data in the country, established and run locally. These requirements are real, but they are narrow and clearly located.

Section 7 set out what could be done, from the steps within easy reach to those that would take years to complete. That agenda is realistic precisely because its foundations are already in place. The markets draw on the same standards, share several regional instruments, and are moving of their own accord towards a common digital rulebook. The convergence so far is partial, and no regional plan produced it. It could also drift apart, because each market can revise its rules on its own. The coordinated steps in Section 7 would build on it deliberately, and make it something banks and supervisors can rely on across borders.

Building that foundation now also readies the region for what comes next. Supervisory and industry attention is beginning to widen from the AI models banks use today to more autonomous systems that can plan and act on their own, and supervisors and the industry have begun to develop safeguards for governing them at the point of action.[144] The risks these systems raise fall on the pillars this report has already traced: dependence on a few providers, cyber exposure, and model risk that travels across markets. The case for common foundations will strengthen as the technology advances.

The people closest to it say the same. One authority observed that **"agents are everywhere now and had completely changed everything,"** while a bank noted that agentic AI needs careful evaluation before it is deployed.

This does not remove the constraints. The convergence is largely voluntary, it operates at the level of principle more often than binding rule, and it reaches its limit where a market requires a bank's systems or data to remain within its borders. Such requirements are becoming more common. The OECD finds that data localisation measures across ASEAN are growing in number and becoming more restrictive.[145] These measures reflect choices each market makes for its own reasons of security and sovereignty, and the report does not understate their effect. They mark the outer limit of a regional approach, and most of the common ground sits well inside that limit. For a region whose banks increasingly operate across the same borders and rely on the same technology, the implication is that much of the groundwork is already in place.

144 Monetary Authority of Singapore with industry participants (BuildFin.ai), Safeguards for Agentic Finance at Runtime (SAFR), industry white paper, version 1.1, 2026

145 Organisation for Economic Co-operation and Development, Digital Trade Review of the Association of Southeast Asian Nations, Section 3.3, 2026. The OECD's count covers measures in all sectors of the economy, including draft measures, across ten member states. It does not apply this report's descriptions of fit.

© 2026 GFTN Limited, All Rights Reserved. Reproduction Prohibited.

09
NEXT STEPS

This report is the first stage of a wider programme of work, led by GFTN and the ASEAN Bankers Association, to build a common baseline for AI governance in the region's banking sector. It has set out where the markets already share common ground, where they differ, and the opportunities that follow.

The next phase of the work will turn those opportunities into a plan. The opportunities differ in what they would achieve and in how hard they would be to reach, so they cannot all be pursued at once. The next phase will therefore work with supervisors and banks across the region to decide which to take forward first. It will then set out an implementation roadmap, describing what each step involves and the order in which the steps could be taken.

The work will also need to align with ASEAN's wider agenda on AI. The ASEAN Working Group on AI Governance, which reports to the region's digital senior officials, leads and coordinates AI initiatives across ASEAN.[146] The programme will seek to ensure that the opportunities identified for banking are aligned with, and complementary to, the working group's work. This would support a coordinated approach to AI governance across the region and avoid duplicated effort. It would also help identify where the two strands of work could collaborate.

The aim throughout is a shared baseline that each market can adopt on its own terms, building on the common ground this report has described.

10
COUNTRY PROFILES

The assessment draws on eleven country profiles, one for each market, each applying the same eight-pillar framework to that market's own instruments and sources. The profiles are held as a separate set, available on request, rather than being reproduced here.



[146] Infocomm Media Development Authority, ASEAN Working Group on AI Governance (WG-AI), accessed October 2026

© 2026 GFTN Limited, All Rights Reserved. Reproduction Prohibited.

11

ANNEX A

Where regional capacity building would help most

This annex looks at the survey results for two groups of markets, to show where shared capacity building would add most value. The groups follow a division that ASEAN uses in its own integration programmes.[147] The first is the ASEAN-6, the six states that were members of ASEAN before 1995: Brunei, Indonesia, Malaysia, the Philippines, Singapore and Thailand. Some economic commentary uses ASEAN-6 for a different set, the five founding members and Vietnam. The second comprises Cambodia, Laos, Myanmar and Vietnam, which joined between 1995 and 1999, together with Timor-Leste. Each group accounts for about half of the banks surveyed: 30 in the first and 31 in the second. Results are reported by group only.

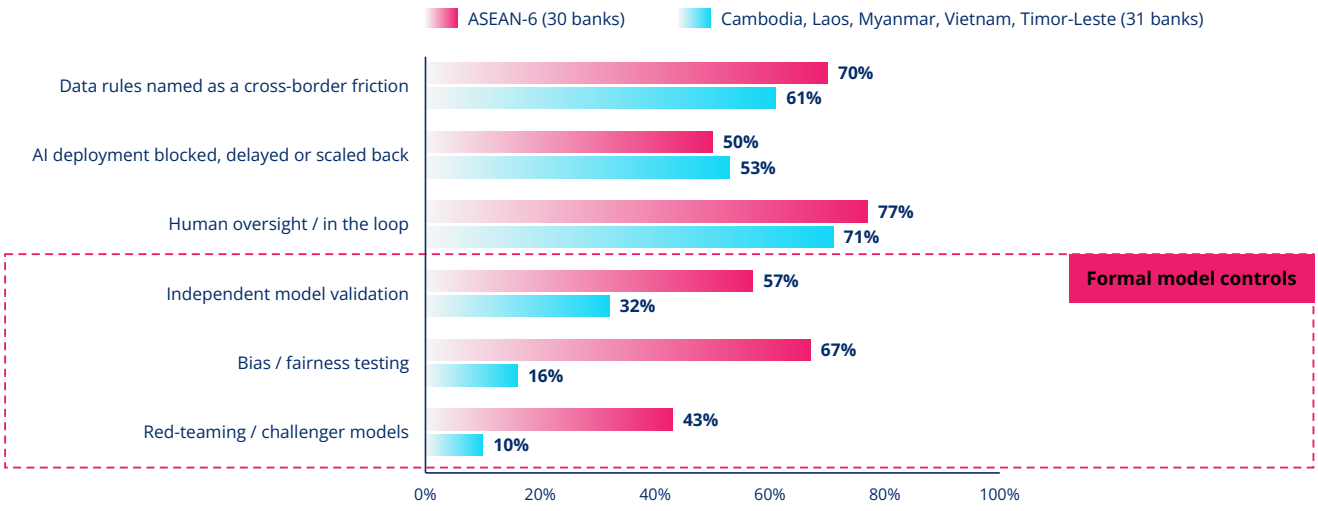
Banks in both groups share the same priorities. Data rules are the cross-border friction they cite most often. Around half the banks in each group have had an AI deployment blocked, delayed or scaled back. Both groups rate consistent governance across markets as highly important, at 4.4 and 4.2 out of 5. Human oversight of higher-risk AI is also widespread in both, practised by more than seven in ten banks.

The groups differ in the formal controls banks have built around their models, such as independent validation, bias testing and red-teaming. Each of these needs specialist skills, tools and data, and banks in the second group report them less often. Two-thirds of ASEAN-6 banks apply bias or fairness testing to higher-risk AI, for example, compared with one in six banks in the second group.

Exhibit 9:

Selected survey results by group of markets

Share of banks in each group giving each answer: 30 banks in the ASEAN-6 and 31 in Cambodia, Laos, Myanmar, Vietnam and Timor-Leste.



Source: the GFTN and ABA survey of ASEAN bank leaders.

This is where regional effort would add most value. Shared testing tools, pooled model-validation expertise and joint training would let banks build these controls without each developing the capability alone. Because banks in both groups share the same priorities, the whole region has an interest in building this capacity together.

The survey was not designed to be representative of each group. Indonesia accounts for 11 of the 30 ASEAN-6 banks, and Vietnam and Cambodia together account for 21 of the 31 banks in the second group. The figures describe the banks that responded and should be read as indicative.

147 ASEAN Secretariat, Initiative for ASEAN Integration (IAI) and Narrowing the Development Gap (NDG), fact sheet, 2021

© 2026 GFTN Limited, All Rights Reserved. Reproduction Prohibited.

12

ANNEX B

How the opportunities trace to the evidence

Exhibit 10 links each of the ten opportunities in Section 7 to the evidence that supports it. The evidence comes from three sources. The first is the survey of 61 banks. The second is the desk research across the eleven markets. The third is the report's discussions with regulators, and with industry participants at the GFTN Select roundtables held across the region. The final column shows which banks each opportunity would help most, based on the report's analysis. Every opportunity would be open to all banks.

Exhibit 10:

The ten opportunities, the evidence behind them and who would benefit most

Opportunity	What the research found	Who would benefit most
1 **A shared vocabulary and governance baseline**	**Survey:** 49% of banks have no internal AI governance policy in force in any of the seven areas of the ASEAN Guide. Banks without cross-border operations have a policy in force in 25% of the areas on average, against 42% for banks with them. **Desk research:** Few markets define AI in binding law, and many rules use the term without defining it. **Regulators and industry:** One supervisor saw value in markets aligning on AI taxonomy and risk management. Industry participants at a GFTN Select roundtable said the industry is invested in using the same taxonomy and language.	All banks, especially those focused on a single market
2 **A shared baseline for transparency to customers**	**Survey:** 69% of banks named transparency and explainability among the three areas most in need of clarity. 34% have an internal policy on it in force. **Desk research:** Seven of the eleven markets add an overlay on transparency, more than on any other pillar. **Regulators and industry:** One supervisor saw value in guidance or practical examples on explainability and transparency. These could show how technical or how plain an explanation should be, depending on whether it is for customers or for supervisors.	All banks
3 **A shared disclosure standard for foundation-model providers**	**Survey:** 39% of banks named model or vendor approval as a source of cross-border friction. **Desk research:** Each bank and each supervisor assesses the same providers separately. **Regulators and industry:** One authority noted that a provider's data-retention terms can change without notice. Industry participants at a GFTN Select roundtable asked how a bank can get a vendor to explain its model.	All banks
4 **Shared tools for testing AI systems**	**Survey:** 67% of banks in the ASEAN-6 apply bias or fairness testing to higher-risk AI, against 16% of banks in other markets. **Regulators and industry:** One authority ran a red-teaming exercise with its banks because global models had not been tested in its language. Local-language prompts triggered jailbreaks more easily than English ones. Another named shared testing methods as an opportunity for regional alignment.	All banks

© 2026 GFTN Limited, All Rights Reserved. Reproduction Prohibited.

Exhibit 10:

The ten opportunities, the evidence behind them and who would benefit most (continued)

Opportunity	What the research found	Who would benefit most
5 **A practical playbook for moving data across borders**	**Survey:** 66% of banks named data localisation and transfer rules as a source of cross-border friction, more than any other factor. This includes 40% of banks without cross-border operations. **Desk research:** In several markets a bank cannot use a standard, pre-approved transfer route. Some markets clear transfers case by case. Others have no settled data-transfer framework yet. **Regulators and industry:** Three regulators raised AI in fraud and anti-money-laundering work without being asked. That work crosses borders by nature.	All banks, most of all those operating across ASEAN
6 **Capacity building**	**Survey:** Banks in three markets described gaps in AI skills, including AI governance and model validation. **Regulators and industry:** Several regulators said they are still building their own AI expertise. Industry participants at a GFTN Select roundtable noted that these capabilities are costly for a single bank to build alone.	All banks, especially those focused on a single market
7 **A shared risk tiering for AI use cases**	**Survey:** 39% of banks named differing risk classifications as a source of cross-border friction. **Desk research:** Markets are starting to set their own risk tiers. One rates banks' AI models as high, medium or low risk. Another's draft AI law proposes four tiers. **Regulators and industry:** Industry participants at a GFTN Select roundtable said most of their AI uses are simple and could be over-regulated without an agreed view of what is high risk.	Mainly banks operating across ASEAN
8 **Portability of model validation**	**Survey:** Banks rate the cost of AI-related compliance across markets at 3.5 out of 5. Banks in three markets said consistent rules could help by cutting duplicated validation or letting one set of evidence serve several regulators. **Desk research:** There is no mechanism for one supervisor to rely on another's assessment of an AI model.	Mainly banks operating across ASEAN
9 **Cross-border sharing of AI incidents**	**Survey:** 57% of banks named security and safety among the areas most in need of clarity. **Desk research:** Every market's incident reporting runs to its own national authority. Deadlines range from one hour in one market to no fixed deadline in another. **Regulators and industry:** One authority noted that a single incident can involve two or three jurisdictions. Industry participants at a GFTN Select roundtable said sharing what has gone wrong would help more than new rules.	All banks
10 **Regional visibility of provider concentration**	**Survey:** 61% of banks rate dependency on a few external model or cloud providers as a high or medium resilience risk. **Desk research:** No market yet looks at provider concentration across banks. **Regulators and industry:** One supervisor warned that a problem at one bank can easily affect others that use the same providers.	The banking system as a whole, and its supervisors

Source: The GFTN and ABA survey of ASEAN bank leaders; GFTN desk research; regulator interviews; GFTN Select roundtables.

© 2026 GFTN Limited, All Rights Reserved. Reproduction Prohibited.

13
SOURCES AND EVIDENCE NOTES

This report reflects the laws, regulations and published guidance in force across the eleven markets as at September 2026, with selected updates to October 2026. AI governance is developing quickly in the region. Several of the measures the report relies on are recent, some enacted only in 2025 and now being brought into effect, and others still in consultation. The findings describe the position as at that date, and the picture will continue to develop as these measures take effect.

The assessment rests on official and primary sources. For each market these include the banking laws, the central bank's notices and circulars, the data-protection laws, the outsourcing rules and any AI-specific instruments. Each is cited in the footnotes where the report relies on it, and set out in full in the country profiles, which are available on request. Each instrument was dated and checked for whether it remains in force, and sources published in a language other than English were consulted in the original. Where an expectation rests on guidance or a measure still in draft rather than on binding law, the report says so at the relevant point.

The report's primary research has two strands. The first is a survey of bank leaders. GFTN and the ASEAN Bankers Association surveyed senior technology, data and risk leaders at 61 banks across all eleven markets between 1 August and 18 September 2026. The survey received 63 responses. One bank returned three, and these are combined so that each bank counts once: a tick-box option counts if any of its respondents selected it, and where they gave different answers to a single-choice question, each answer carries an equal share. Figures are expressed as a share of the banks that answered the question concerned, so the base varies slightly between questions.

The average rating of 4.3 out of 5 for consistent governance, for example, comes from the 48 banks that gave a rating. The responses are self-reported, and they are intended to show what the region's banks are experiencing rather than to form a statistically representative sample. Because the number of responding banks varies across the eleven markets, the survey is read for region-wide patterns rather than to compare one market with another. Vietnam, Indonesia and Cambodia together account for 52% of responding banks. To check that this does not drive the results, each headline figure was recalculated with every market weighted equally. The findings hold. The largest shifts are in generative AI in production, from 51% to 64%, and in accountability as an area needing clarity, from 51% to 61%. Neither changes a conclusion in the report. Annex A reports results for two broad groups of markets, never for a single market. No response is attributed to an individual bank.

The second strand is a set of interviews with the region's supervisors, together with GFTN Select roundtable discussions. These were held under the Chatham House Rule, so the report draws on what was said without attributing it to any individual or institution. Contributions from supervisors appear as "a regulator", "a supervisor" or "an authority", contributions from banks as "a bank", and contributions from the roundtables as "industry participants". Quotations are lightly edited for readability, without changing their meaning.

The report describes how far the markets' rules work together across borders. It does not measure how advanced or how effective each market's framework is, and it does not rank or score the markets. Its descriptions of common ground, overlay and local implementation are the report's own reading of how each market's rules fit with those of its neighbours, drawn from the underlying instruments rather than from any label a market applies to itself.

14
GLOSSARY

This glossary explains technical terms used in the report.

Agentic AI. AI systems that can plan and carry out a sequence of actions with limited human direction.

Disparate-impact test. A check of whether a model's outcomes fall unevenly on different groups of customers, even where the model does not use a protected attribute directly.

Drift. A decline in a model's performance over time, as the data it sees in use moves away from the data it was trained on.

Foundation model. A large AI model trained on broad data that can be adapted to many tasks. The large language models behind generative AI are examples.

Herding. Many institutions acting in the same way at the same time because they rely on similar models or data. This can amplify market moves.

Jailbreaking. Writing inputs that get an AI model to ignore its safety rules.

Red-teaming. Testing a system by deliberately attacking it, as an adversary would, to find weaknesses before others do.

© 2026 GFTN Limited, All Rights Reserved. Reproduction Prohibited.

GLOBAL FINANCE & TECHNOLOGY NETWORK (GFTN)

6 Battery Road, #28-01, Singapore 049909
gftn.co | hello@gftn.com

This document is published by Global Finance & Technology Network Limited (GFTN) as part of a programme of work led by GFTN and the ASEAN Bankers Association. The findings, interpretations, and conclusions presented in this report reflect the views of the author(s) and do not necessarily represent those of GFTN, its Board, management, stakeholders, or any individual participant and their respective organisations.

© 2026 Global Finance & Technology Network Limited, All Rights Reserved. Reproduction Prohibited.