



AI in Indian banking: Transformation through trust and technology

January 2026



Foreword



Gayathri Parthasarathy

India Financial Services Leader and
Global Financial Services Technology Leader
PwC India

Over the past few years, Indian banking has witnessed remarkable changes—from Unified Payments Interface (UPI) reshaping payments to mobile wallets and online banking platforms that provide access to banking services anytime, anywhere. Anchored in personalisation, speed, and convenience, the digital wave is changing the way banks interact with customers.

Poised to lead the next wave, AI is set to accelerate transformations across customer engagement, risk management, and banking operations. PwC's 28th Annual Global CEO Survey revealed that nearly half the respondents expect GenAI to increase profitability in the year ahead, yet only one-third have AI embedded in key processes.¹ Our conversations with banking leaders reveal a sector torn between ambition and action. While pilots are abundant, there is hesitation to commit to enterprise-wide adoption.

This study focuses on this hesitation between pilots and adoption. Our research shows that trust lies at the crux of this hesitation. While AI's potential is widely accepted, trust in how this technology will be governed and regulated, implications of its failure on reputation and customer trust, trust in longevity of invested infrastructure, and banks' maturity to incorporate responsible AI are holding the sector back.

The Indian regulatory landscape is constantly evolving. Governing bodies and regulators are still figuring out the mechanics of effective governance. RBI's Framework for Responsible and Ethical Enablement of Artificial Intelligence (FREE-AI) and the Digital Personal Data Protection (DPDP) Act are starting points. The Government-backed AI mission and indigenous platforms have opened up options for banks. The building blocks for adoption are in place. What remains is the bridge between intent and implementation.

This report offers that bridge. We have structured our findings and recommendations around the theme of trust-led transformation for AI in Indian banking. This study elaborates on frameworks for governance, infrastructure selection, model testing, and strategies for short-long term AI adoption. We have provided actionable guidance for institutions ready to move beyond pilots.

We believe that the opportunity is here. We have access to the right technology resources, and the timing is right. What is needed is a leap of faith from Indian banks. We hope that this report serves as a practical guide for banks seeking to unlock AI's transformative potential while managing its inherent risks.

1 <https://www.pwc.com/gx/en/ceo-survey/2025/28th-ceo-survey.pdf>

“



Across the banking sector, 65% of institutions have begun adopting AI, yet only 9% have reached production-grade deployment. Our mandate is clear: Convert pilots into trusted, explainable, and human-in-the-loop solutions that scale without compromising customer confidence.”

Vivek Belgavi

Partner and Leader, Financial Services Technology
PwC India

“



Modern AI has pivoted from reliance on structured curated data and exponentially broadened its scope of insights and intelligence by consuming unstructured and semi-structured data. AI has changed the game—redefining workflows, inspiring new data models, and unlocking holistic insights. However, meaningful adoption relies on strong governance and alignment with regulatory expectations.”

Hetal Shah

Partner, Data and Analytics
PwC India

“



Banks are right to be cautious about the risks of AI. However, this should not hinder them from integrating AI in value chains. Banks should embed robust AI risk frameworks, building confidence in responsible innovation while maintaining accountability.”

Anit Shanker

Partner, Financial Services
PwC India

Table of contents

Executive summary	7
01 Introduction	8
02 Methodology	10
03 Key findings and recommendations: Indian banking sector	12
04 AI democratisation framework	30
05 AI adoption in India: Journey so far	32
06 Conclusion	35





Executive summary

The path to AI adoption is not a simple one, and many Indian banks appear to get caught at the use-case or pilot stage. Some studies have reported that organisations working on a few strategic bets have a higher success rate compared to those running small pilots in many areas.² One of the key reasons for AI programmes not taking off as intended is lack of trust in these systems. Trust can be defined as a combination of accuracy, consistent output, observability, explainability, and data privacy in the context of organisations, especially banks and regulated entities (REs). AI is not a lift and shift technology that can be ingested like software. It requires a well thought out approach on aspects such as infrastructure, models, observability, organisational setup, and AI literacy. In the past few years, we have seen both government and regulators focus on these aspects of AI enablement.

This report analyses the maturity and adoption of AI within the Indian banking sector, drawing insights from primary (survey and interviews across banks and financial institutions) as well as secondary research. It delves into the sector's concerns about gaps in the country's AI policies. We've attempted to find answers to these concerns, which we believe are holding organisations back from truly realising the transformational potential of this technology. We have provided specific recommendations on risk and governance, infrastructure and architecture setup, model lifecycle, and AI strategy and investments. These recommendations have been arrived at through successful cases of AI adoption within the banking industry, both in India and globally.

The report also presents multiple frameworks and approaches that can be adopted by REs to navigate the adoption journey. We've also included a template followed by the Government of India for building the foundational elements of AI adoption. Regulators and industry bodies can follow a similar path to bring together the ecosystem and enable trusted adoption of this technology for transformation.

2 <https://fortune.com/2025/08/18/mit-report-95-percent-generative-ai-pilots-at-companies-failing-cfo/>

01

Introduction

Changing dynamics in the banking industry

India's AI market was valued at \$9.51 billion in 2024 and is projected to grow from \$13.05 billion in 2025 to \$130.63 billion by 2032, exhibiting a CAGR of 39% over the forecast period.³

These forecasts suggest that AI will play a key role in shaping the next 5-7 years and beyond. By using AI to unlock transformation opportunities across operations, compliance, risk assessment, and customer experience, the banking industry can emerge as a major driver of this growth.

Customer expectations: Emphasis on convenience

There has been a noticeable shift in customers' banking behaviour over the last few years, with customers beginning to value the convenience offered by digital banking. While physical branches continue to remain important for building trust and offering complex products (loans, wealth management), digital modes are now preferred for routine transactions.

This is evident from the 38x surge in the volume of digital payments (CAGR of 46%) in India over the past decade. During the same period, the value of digital transactions has tripled, rising at a CAGR of 10%.⁴

Banks that integrate AI in the customer journey to facilitate instant decision-making and enable personalised customer engagement are bound to be the front runners. A few examples are AI-powered fraud detection, AI-integrated credit decisioning and straight through processing, AI-powered personalised nudges to relationship managers, and AI for facial recognition and spoofing detection for smart customer e-KYC.

Tech expectations: Embedded 'AI workforce'

Evolving regulations and disruptive AI capabilities like GenAI and agentic AI are redefining risk identification, efficiency, compliance, and customer experience. The concept of an AI workforce is no longer futuristic; it's slowly becoming a mainstream reality. What started with small pilots and support chatbots deployments has gradually matured to enterprise-level initiatives.

Access to large and complex volumes of data and introduction of large language models (LLMs) and agent-based automation has pushed banks to rethink AI adoption in customer interactions, product innovation, operations, and risk and compliance. Specifically in risk, with improved AI technology, some global banks are becoming increasingly better at fraud identification, anti-money laundering, payments, etc., thereby making the banking industry more robust.

Trustworthy AI enablement

The rapid pace of AI maturity and increasing technology savviness of customers is posing challenges to banks and regulators today, leading them to rethink their AI and data strategy and keep pace with innovations by redesigning their operating models.

In the last few years, we have seen strong encouragement for banks to embed AI in value chains. Enhanced Access & Service Excellence (EASE) mandates (EASE 6.0, EASE 7.0, EASE 8.0) have been outlined and released annually by the Department of Financial Services⁵ under the Ministry of Finance to improve the performance of public sector banks (PSBs). This, along with the RBI's Framework for Responsible and Ethical Enablement of Artificial Intelligence (FREE-AI) report,⁶ is a key pillar for Indian banks outlining their AI adoption strategy.


3 <https://www.fortunebusinessinsights.com/india-artificial-intelligence-market-113969>

4 <https://rbi.org.in/scripts/PublicationsView.aspx?id=23436>

5 https://www.iba.org.in/reports/ease-8-0-launch-report_1703.html

6 <https://rbidocs.rbi.org.in/rdocs/PublicationReport/Pdfs/FREEAIR130820250A24FF2D4578453F824C72ED9F5D5851.PDF>

Evolution of EASE mandates



Examples of AI implementation in the banking sector (India and global)

- A large PSB has built an AI chatbot for improved customer service and query resolution.
- A large private sector bank has embedded AI for voice biometrics and ML risk scoring, enhancing security and conversion rates.
- Large multinational banking institutions have leveraged natural language processing to automate contract review and compress hours of manual legal work into seconds while strengthening compliance oversight and freeing legal teams for higher-value strategic work.
- A large multinational payment services institution has used AI platforms to power real-time fraud detection across billions of transactions, cutting false positives and improving detection rates.



02

Methodology

This study employed a dual-lens methodology combining inside-out expertise with outside-in primary research.

The **inside-out approach** tapped PwC’s sectoral expertise and first-hand experience in AI strategy, governance, and implementation, enabling interpretation of findings within practical industry realities. This report is grounded by secondary research on global benchmarks and regulatory frameworks such as RBI’s FREE-AI guidelines, the UK Financial Conduct Authority’s AI adoption report (2024),⁷ and the Digital Personal Data Protection (DPDP) Act.⁸


The **outside-in approach** drew on primary research facilitated through a curated survey that captures the state of AI/GenAI adoption across forty-three respondents from the Indian banking and financial services ecosystem. The survey covered two broad areas: (1) the business/functional aspects of AI such as adoption, usage, value, and talent; and (2) the technology, compliance, risk, and governance aspects of AI.

We also combined the insights from our survey with those from interviews with a few subject matter experts (SMEs) on AI. Due to the limited sample size, some sample bias might have come in, especially owing to the larger proportion of private banks. Nevertheless, the survey and interviews together provide us with a ‘directional’ view of the challenges faced by the industry.

Respondent overview:

The respondents represent a diverse mix of private sector banks, large PSBs, small finance banks, non-banking financial companies (NBFCs), foreign banks, and a FinTech entity, with a strong tilt towards larger balance sheets: 33.3% of entities report assets above ₹5 lakh crore, while a further 26.7% report assets between ₹1–5 lakh crore. Private sector banks dominate the respondent pool at 50%, followed by large PSBs at 20%, and smaller institutions and FinTechs collectively accounting for the remaining 30%.


Distribution of respondents by asset size



7 <https://www.bankofengland.co.uk/report/2024/artificial-intelligence-in-uk-financial-services-2024>

8 <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2190655®=3&lang=2>

Distribution of respondents by level of seniority



Survey insights: Respondents hold mid-to senior-level positions with direct influence over AI strategy and implementation. Senior leadership and head/director roles together account for 33.3% of respondents, while VP level and senior manager positions contribute a further 28.6%. Executive-level participation stands at 9.5%, with manager and analyst/staff roles collectively representing 28.6%. This profile mix captures perspectives from practitioners who are both close to strategic decision-making and hands-on implementation of AI/GenAI initiatives.

Design and process:

- **Question types:** Multiple-choice with rating scales (1 to 5), and open-ended questions to capture both quantitative benchmarks and qualitative insights
- **Mode:** Online questionnaires supplemented by interviews for deeper context
- **Validation:** Global benchmarks and regulatory frameworks were used to corroborate and analyse the findings and present recommendations



03

Key findings and recommendations: Indian banking sector

Leveraging PwC's consulting experience in the banking sector, our survey findings as well as the interviews conducted with SMEs working in the field of AI/GenAI, we have collated key observations on AI adoption challenges faced by Indian banks.

We conclude that there is no single correct or incorrect pathway towards AI adoption at scale—each bank needs to undertake its own journey depending on its starting maturity,

tech-stack availability, risk appetite as well as aspirations towards becoming an industry pioneer in AI adoption, being a risk-averse follower, or landing up somewhere in the middle.


We will address three aspects from an AI adoption perspective—where Indian banks stand today, AI application appetite of the banks, and hindrances to AI adoption.



3.1 Adoption maturity—where Indian banks stand today

The Indian banking sector’s journey could be described as one in active transition, with widespread pilots/experimentation but limited enterprise-scale deployment.

Overview of institution-wise AI adoption maturity



Survey insights: About 65% of respondents indicated they have begun adopting AI. Although running of pilot projects is common across the banking sector, only 9% (comprising private institutions and small finance banks) have reported active production-level applications. Most institutions have adopted a cautious approach with near-term and short-term planning for AI use cases.




Our analysis: Most private banks are proactively looking to prioritise AI in their strategic vision. They are trying to get a head start on embedding AI in their value chain from applications in business models such as customer interaction, re-engineering workflows, and leveraging data for smarter decision-making. While PSBs continue to navigate the ‘trust gap—from a governance, risk, and ROI perspective—they have very closely aligned themselves to the Ministry of Finance’s EASE mandates of digital-first banking and are now moving towards AI-led banking.

3.2 Indian banks' application appetite

A deep dive into Indian banking's AI investments suggests that banks are winning the efficiency battle but losing the transformation war.

Institution-wise AI investment priorities



Survey insights: AI investment priorities reveal a defensive posture. Operational efficiency and risk management/compliance together command 50% of all investments. On the other hand, new product development such as trade finance, personalised banking, and investment advisory collectively account for less than 3%. The sector is optimising for cost and compliance today, but under-investing in the innovation foundations needed for tomorrow.



Our analysis: Indian banks are investing in AI, and their cautious approach towards its application is evident in their use case focus. Across institution types, AI is overwhelmingly deployed to streamline operations, enhance customer experience, and tighten risk controls—areas where short-term ROI feels tangible and safe.

This reveals a gap in the adoption strategy, where AI is powering operational excellence, but its potential to fuel bold growth strategies remains untapped. In a market where digital-savvy competitors are revamping customer engagement, Indian banks must pivot from defensive automation to offensive innovation—using AI not just to

optimise existing processes but to create new products, unlock new markets, and redefine customer value. The next frontier isn't about doing the same things faster; it's about doing fundamentally different things. Until that shift happens, AI will remain a powerful engine for cost control, but not the catalyst for competitive advantage.


3.3 A closer look at what's hindering AI adoption in Indian banking

In a world that's being rapidly transformed by AI, three factors define success in AI adoption: trust, technology and transformation. This section delves into each of these factors in the context of Indian banking.

3.3.1. The trust story: Addressing risk and governance concerns

While the RBI's FREE-AI report and the DPDP Act together provide a strong foundation for trustworthy AI, banks today are still figuring out how to operationalise these recommendations into model lifecycle controls and third-party oversight.

Overview of constraints in Indian banking



Survey insights: On the **regulatory** front, know your customer (KYC) and anti-money laundering (AML) compliance emerges as the most formidable barrier. Similarly, consumer protection regulations and the DPDP Act are shaping deployment strategies. Banks perceive regulatory ambiguity as the primary obstacle, with 80% citing lack of clear guidelines as a challenge.

On the **non-regulatory** front, concerns around customer acceptance, legacy system integration, and budget constraints are common roadblocks that impact AI adoption.




Our analysis: Indian banks are grappling with the challenge of turning AI ambitions into practical and regulatory- and risk-compliant operations. Most Indian banks are stuck in pilots, largely because the true picture of ROI and risk visibility in full-scale production is yet to completely come to light. The biggest hurdle lies in translating policy-level expectations into everyday controls within business-as-usual processes.

While most banks and enterprises are looking to develop AI solutions in-house, a broad spectrum of third-party AI solutions is already touching every aspect of day-to-day work. In the past few months, the world has seen a variety of threats emerge, specifically in the case of apps which embed AI and GenAI—for example, prompt injections, data poisoning, echo leaks, and cypher attacks. These cyberthreats can also lead to reputational risks, data leaks, identity theft, sensitive information leaks, app shutdowns, and widespread fraud.

Indian banks fear that a single AI failure could severely damage customer trust and reputation.

Risk themes prioritised by Indian banks



Survey insights: Reputational risk emerges as the biggest AI-related concern, with a score of 4.33. It is followed closely by third-party dependencies at 4.05. Conversely, model bias and fairness 3.11—which is a foundational attribute of responsible AI—lie at the other end of the spectrum.



Our analysis: Without enterprise-wide AI governance (beyond the chief technology officer [CTO]), banks risk reputational harm and regulatory friction. Governance maturity must catch up with adoption. The challenge, as we see it, is not limited to technical feasibility but also ensuring AI buy-in from leadership, process/business owners, frontline users/customers as well as regulators and other key stakeholders.


Governance should be an enterprise-wide agenda and not the sole responsibility of the technology team.

While the RBI has developed a framework for responsible AI, it's still not mandatory for banks to implement it. Further, some banks remain unsure of the kind of controls they must embed with respect to AI governance, and they are waiting for regulators to provide more clarity before they start on their AI journey.

In the near future, a standardised approach for quantifying value generated vis-à-vis risk taken is needed. The drivers for value should not be only ROI, but also improvement in the

bank's asset quality, customer experience, and robustness in operations. On the risk side, data privacy, reputational impact, systemic shutdown, AI accuracy, fairness and explainability, and operational viability should also be taken into consideration.

Banks that successfully convert AI policies from a sectoral mandate to a growth enabler and risk-resilient capability will reap the highest benefits and lead the next wave of transformation.



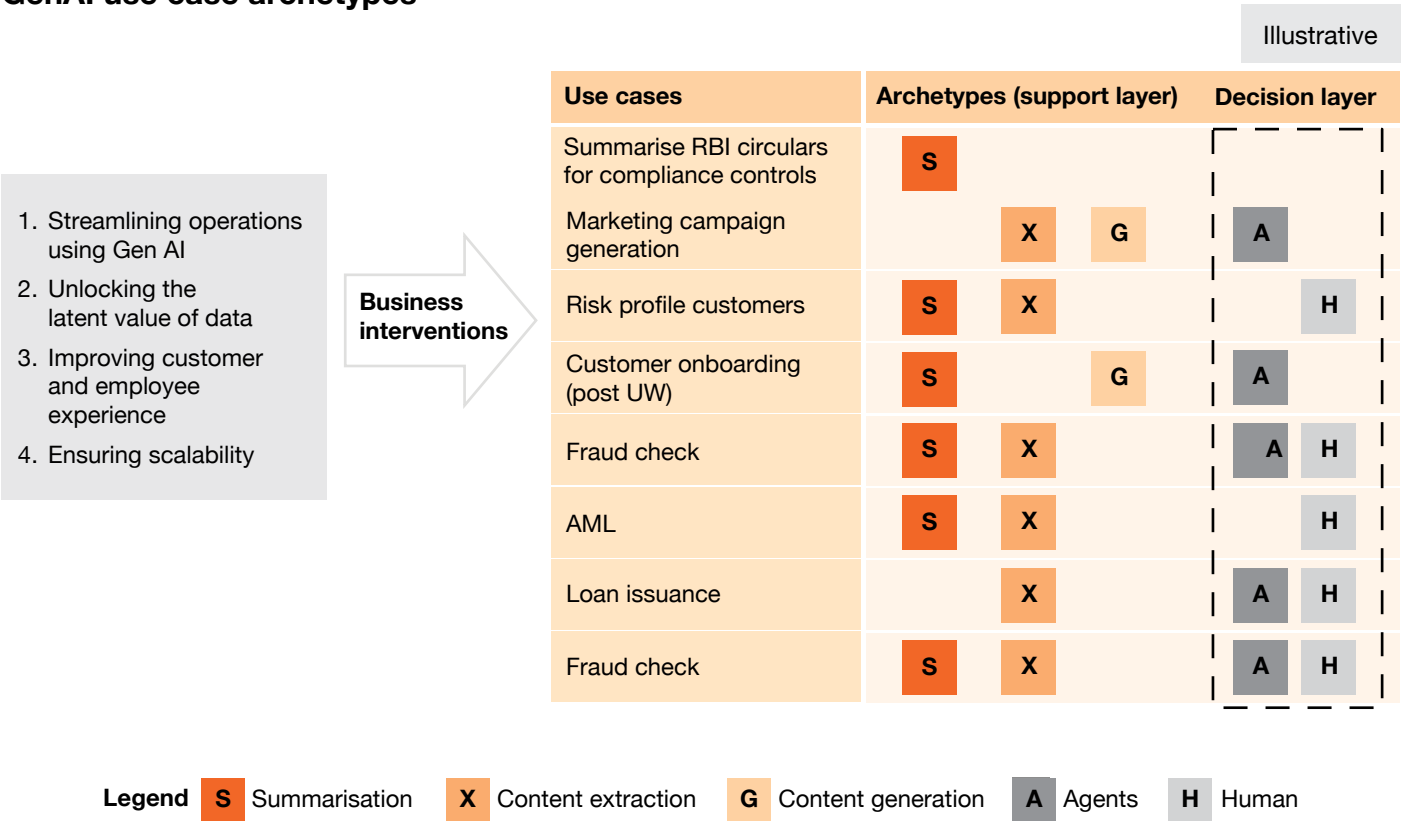
Our recommendations

Indian banks have outlined valid concerns over AI failure as ungoverned AI and GenAI can cause significant disruptions in banking operations, reputation, and customer trust. We have outlined the best practices that would help mitigate these risks.

1. Kick off with an AI assistant rather than autonomous interventions

While banks wait for a mature governance setup, in the interim, they can be strategic about the kind of AI interventions they invest in. Typical AI use cases would be a combination of one or all components of the archetypes defined below—summarisation (S), info/content extraction (X), content generation (G), AI decision-making agent (A) for trivial decisions, and a human in the loop (H) for specific and strategic decision-making.

GenAI use case archetypes




The figure above outlines samples of AI use cases. Simple AI use cases use one or two AI capabilities with an agent or human oversight (e.g. summarising RBI circulars, generating

campaigns). Complex ones involve multiple autonomous AI interventions with built-in guardrails for edge cases, fail-safe mechanisms, and human in the loop.



Six stages of human–AI collaboration



The above framework outlines how integration and interactions between human and AI systems would evolve over time with AI adoption maturity. The journey from ‘AI as advisor’ (citizen-led productivity chatbots) to ‘AI as a decision maker’ or ‘AI as a self-learner’ would entail building stricter guardrails for AI governance, embedding risk scoring for learning systems, bounded autonomy (maker-checker for high-risk activities, exit loops), explainability and traceability, observability, robust testing frameworks, and a responsible, accountable, consulted, and informed (RACI) matrix.


We suggest banks prioritise applications of ‘AI as assistants or co-creators’ where human oversight is integral, minimising risk and building trust in AI. Post achieving stability and credibility in the output from AI assistants, banks can look to expand to agentic AI use cases that are more complex and require a higher level of autonomy.

2. Build a multi-tiered governance framework (structural concept)

AI adoption will not be without risk. As in the case of any innovative technology, we should assume that there will be mistakes from systems, shortcomings in the governance processes, and errors in judgement by humans. We need to address these through a clear mitigation plan that pivots around cybersecurity and data protection. The plan should enable cross-functional collaboration (internal) and participation from all relevant stakeholders (external). It should be based on the guiding principles of transparency, agility, flexibility, and correctness. The presence of such a governance framework will go a long way in enabling AI to evolve safely.

We recommend that banks build multi-tiered governance—that is, lines of defence that safeguard each step of the AI lifecycle from development and deployment to usage and maintenance, and establish clear accountability with clearly outlined responsibilities.

Multi-tiered AI governance framework



For instance, your first line of defence would help banks address various facets of risk and controls:

- At the development stage—identify any malicious data or libraries that are being used by the developers
- At the deployment stage—cater to red-teaming activities before the launch
- At the usage stage—safeguard prompt, defining permissible inputs and outputs

The second line of defence, with an AI review board made up of representatives from the compliance and legal, data and privacy, cyber, and IT teams and the AI COE, would help banks in the following manner:

At the development stage:

- The AI COE would provide AI sandboxes for experimentation along with approved reference architecture.
- The risk team would outline guardrails around data lineage and protection, establish domain-specific guardrails around lending, financial due diligence, customer profiling, etc.

At the deployment stage:

- The AI COE would help maintain model registry—versioning, metadata, and controlled releases—and ensure capacity planning.
- The risk team would validate operational readiness for production release based on model and data drift, fail-safe measures and escalation matrix, and privacy and security aspects such as response filters and incident mitigation playbooks.

At the usage stage:

- The AI COE would undertake centralised monitoring of use cases in production and during the model lifecycle (including decisions on archiving/ sunsets).
- The risk team would review auditable logs of compliance controls and mitigations, ensure continuous monitoring, and enforce RACI compliance.

Similarly, the third line of defence would strengthen governance at the enterprise level:

- Establish overall organisational strategy and principles for leveraging AI.
- Align and maintain communication with regulators and other stakeholders.
- Guide organisational direction and priorities for governance and ethics.
- Define risk tolerance for the organisation and business units.
- Address issues escalated by the AI review board.

3. Bake in observability into AI systems through a governance platform (orchestration layer)


While banks initiate AI pilots with human oversight, the next step would be to embed a robust observability layer in the AI system. This is an essential step to build trust, compliance, and traceability in a highly regulated banking setup.

An ideal observability layer would provide traceability for every decision, interaction, and outcomes across AI systems—from data, models, agents, and workflows to model outputs and user feedback.

The key capabilities of this layer would cover transparency from the following aspects:

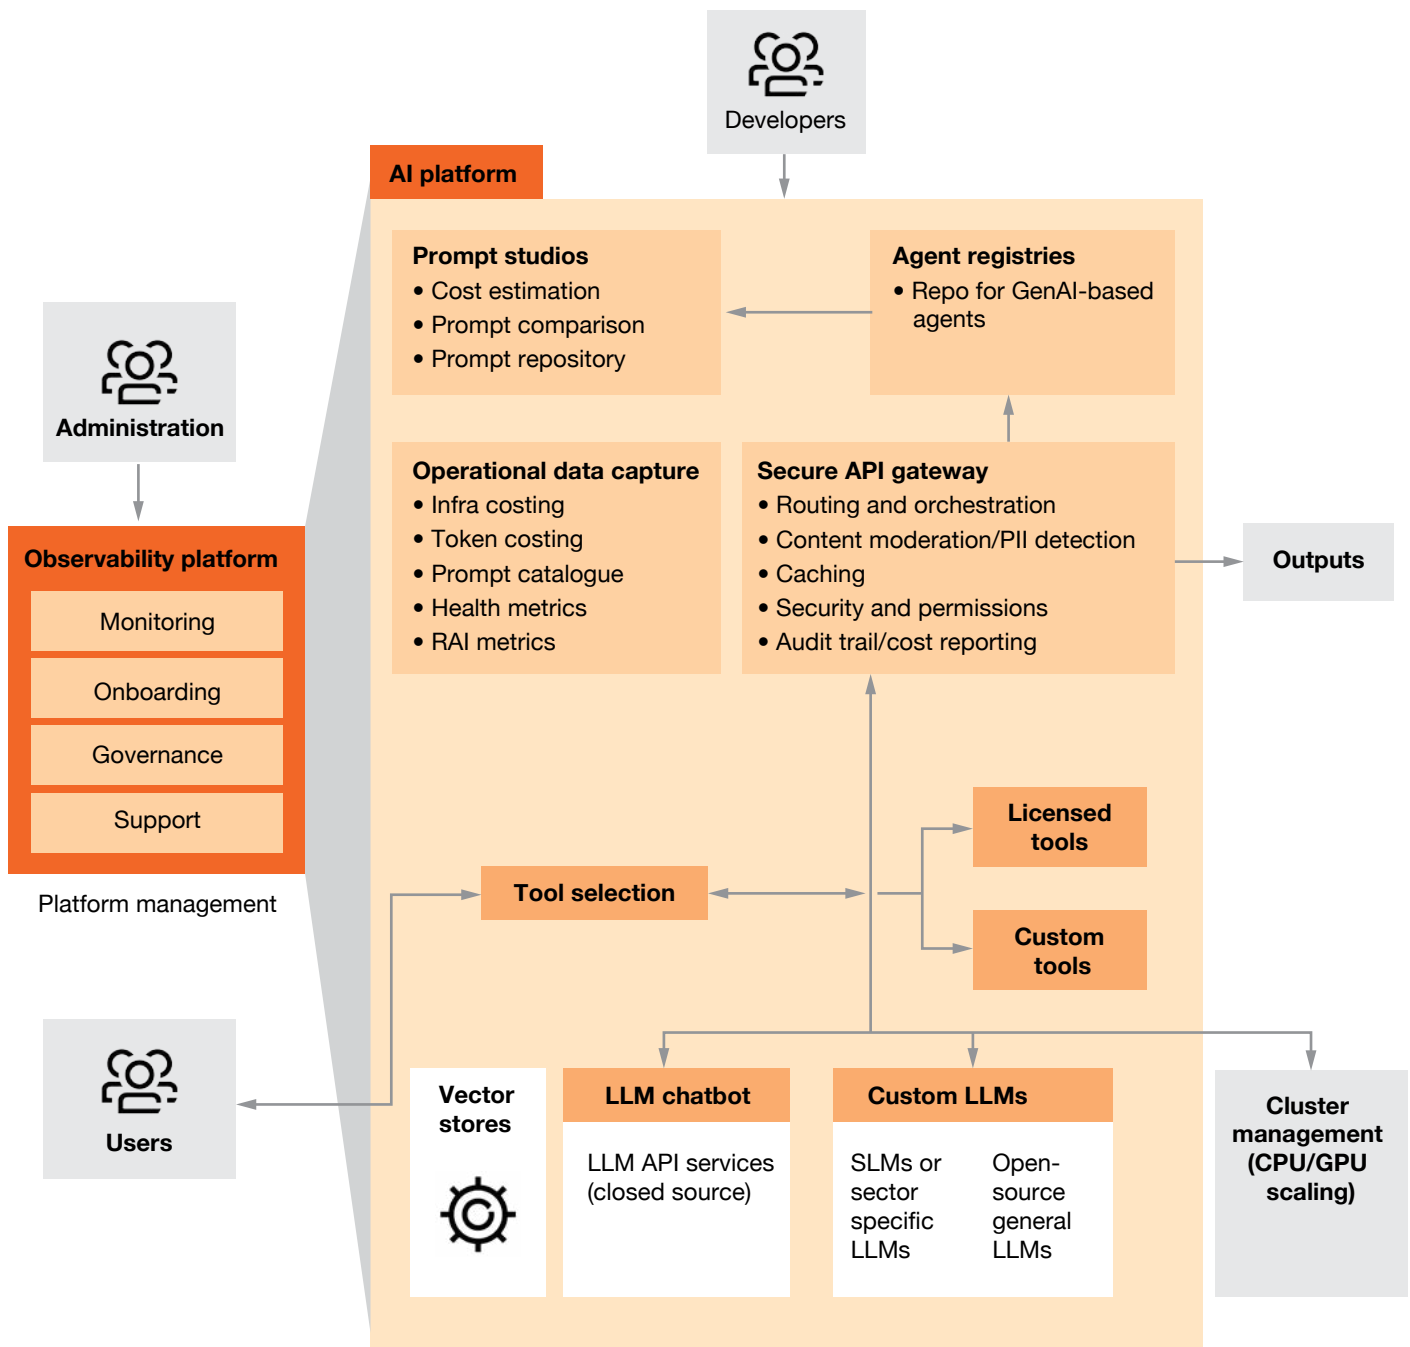
- **Agent catalogue and management:** Effectively catalogues AI agents and workflows with visibility into lifecycle, logs, performance, actions, and health.
- **Vector database catalogue:** Monitors knowledge bases and embeddings for storage, performance, and integrations.
- **Model catalogue:** Functions as a central registry for all LLMs and models with version controls, performance metrics, lineage, and dependency tracking.
- **Event catalogue:** Helps maintain an audit trail for compliance by tracking events and anomalies.
- **Dashboard and analytics:** Offers real-time and historical insights with custom reports.
- **Governance and compliance:** Enforces data policies and regulatory controls.

Components of GenAI observability layer



Banks should adopt a **platform-based approach** to ensure governance, observability, and accessibility. This platform would act as an orchestration layer bringing together developers, tools, LLMs, databases, and users along with governance individuals.

Observability platform setup



The platform would ensure:

- **Unified monitoring:** Visibility across data, APIs, models, incident reports, and outputs
- **Security and compliance:** Role-based access, personally identifiable information (PII) protection, audit logging, and regulatory compliance
- **Scalability:** Seamless orchestration and resourcing for licensed and custom tools

Observability platform in action:

An overview of how the platform would monitor an AI use case:

Case study: Loan approval automation

- **The model catalogue** registers the LLMs used for credit scoring model (versions, dependencies, and performance metrics).
- **The event catalogue** logs every decision made by the model, including risk scores and approval/rejection outcomes.

- **Dashboard and analytics** provide real-time visibility into approval trends, highlighting anomalies such as sudden spikes in high-risk approvals.
- **Governance and compliance** ensures that RBI guidelines on fair lending are enforced by monitoring bias metrics.

This example outlines how different components of the platform work together to deliver end-to-end visibility, continuous monitoring, and trust in AI for banking processes.

Key takeaways:

- Establish a board-approved AI policy.
- Build SOPs for AI controls—RACI, explainability, threat/incident/breach reporting, and third-party model audits.
- Templatised model lifecycle using the FREE-AI sutras and embedding DPDP compliance by design into every AI workload.
- Enable human autonomy over AI decisioning for sensitive and critical areas.




3.3.2 The technology story

3.3.2.1 AI adoption in India – scaling infrastructure

Banks are no longer confined to a one-size-fits-all approach for AI infrastructure setup; instead, they are thinking outside the box. They are willing to explore hybrid setups that feature infrastructure aligned with data sensitivity, compliance, and business priorities.

Infrastructure choice by institution



Survey insights: Cloud infrastructure dominates AI/ML deployments across Indian financial institutions, with 71% of respondents choosing cloud solutions, led predominantly by private sector banks who account for 55% of total adoption. On the other hand, on-premise graphics processing unit (GPU) setups remain marginal at just 13%, signalling a decisive industry-wide shift toward scalable, hyperscaler-driven AI infrastructure.

A sizable chunk of AI investment goes into the tech setup which enables rapid prototyping and deployment. Most banks are still contemplating to whether they should invest in a cloud-first or on-premise setup, viable options for the long run, and how they can ensure data protection and localisation imperatives under the DPDP Act and sectoral guidance.

While neither approach is right nor wrong, banks should weigh the following considerations while choosing either of the two, or a combination of both.

Hosting considerations for AI adoption

Considerations	Cloud	On-premise
Speed vs control	Rapid prototyping, development and deployment capabilities	Greater control on infra, sensitive data, compliance and governance
Adaptability	Access to latest innovative AI capabilities	Focused on stability and slower access to latest AI capabilities
Cost	Pay-as-you-go models with rigorous governance on usage and spends	Procuring infra is capital and time-consuming investment
Maintenance constraints	Infra would be external vendors lookout, but internal usage must be governed	Requires specialised infra management skills
Cybersecurity and vendor risk	Third-party risks and potential exposure to breaches	Reduces external dependencies but focus on internal protocols

Our recommendations

1. A strategic approach to scale enterprise technology

Banks that are eager to embark on their AI journey need not wait for the enterprise to be set up. They can incorporate the hybrid approach of starting with low-risk use cases in cloud, and eventually graduating sensitive workloads to on-premise/sovereign environments as getting the right infra in place can take time.

Banks can prioritise use cases which are low use, internal facing, and do not require sensitive data or a dedicated governance layer for cloud-based hosting.

Given below are six decision levers for data sensitivity, ROI, scale, speed, customisation and control, and legacy system integrations that could help banks evaluate their hosting strategy.

Some of the parameters which banks should consider during the selection process are:

- Cloud:** Data sensitivity is low, the solution is internal to the bank, speed to market matters, volumes for data and usage fluctuate, and the setup needs to be convenient for rapid prototyping and testing.
- On-premise:** Data sensitivity is high, models need to be customised, the solution needs to be integrated with legacy systems, and governance is a prerequisite.
- Hybrid (cloud + on-premise):** Data sensitivity is moderate, use cases are customer facing, and gradual growth in volumes is expected for data and usage.

Framework for hosting strategy based on use case maturity

Decision levers							Complexity score from 1 to 5
Use cases	Data sensitivity and compliance	ROI and cost efficiency	Scalability and usability	Speed to roll-out	Model customisation and control	Integration with legacy systems	
Sales channel coaching agent	2	4	5	5	3	3	Cloud
Customer-facing chatbot	3	5	5	5	3	4	Cloud/ on-prem
Credit memo generation with risk profiling	5	3	3	3	5	5	On-prem

Key takeaways:

- A hybrid approach is the way forward, and banks should adopt a strategy that is not strictly tied to any one type of infrastructure.
- Banks should be agile and plan for technology obsolescence, so they can pivot to newer tech.
- Inaction or long procurement cycles could be costly, and the first movers will make a substantial impact in the market share of others.




3.3.2.2 The reliability dilemma: Adopting smarter testing strategies

Manual QAs and security clearances could lead to delays in production deployment cycles.

AI's reliability is non-negotiable for the Indian banking sector, as banks have to operate under strict RBI guidelines, DPDP Act mandates, and sectoral compliance norms such as KYC, AML, and consumer protection

Though AI testing and validation are extremely critical steps in the model lifecycle, they often lead to deployment bottlenecks. AI failure can trigger penalties, reputational damage, and affect customer trust, and this underscores the importance of testing.

Pilots are often stalled before production due to inefficient testing strategies that are highly reliant on manual testing or traditional QA. As a result, they face roadblocks in UAT and information security clearance. Current testing methodologies do not ensure completeness and consistency as they don't adequately demonstrate resilience against data leakage, prompt injection, or compliance breaches. Additionally, AI language models lack explainability due to the underlying neural network architecture.



Our recommendations

1. Augment testing methodologies:

Leveraging hybrid red-teaming methodologies that combine human evaluation with a 'judge' LLM scoring can stress-test fairness, and simulate scenarios for compliance and robustness at scale to minimise instances of unsafe responses, unintended actions by AI, and data leaks. It covers holistic scenarios of test cases far beyond manual capacity.

Another approach could be triangulation of outputs across multiple models and cross-verifying outputs from the primary LLM, secondary LLM, and domain-specific SLM to ensure reduction in single model failure risk.

2. Client zero approach:

A client zero approach follows the principle of testing solutions in a near-production setting for internal teams. Banks can leverage the fact that bank employees are themselves customers of the same bank. A planned release of AI solutions for these employees could capture valuable feedback, check for any edge cases in a near-production environment, and enhance guardrails before rolling out solutions to customers at large.

Key takeaways:

- With AI systems becoming more sophisticated, scaling them with human testers at the same pace will pose challenges. Use AI to review AI.
- Leverage LLM evaluation frameworks such as HELM, RAGAs, and DeepEval for LLM selection and output evaluation.

3.3.3. The transformation story: The redundant AI investments trap

There are growing concerns about today's AI investments becoming obsolete as emerging AI technology and models continue to get faster, cheaper, and smarter at a rapid pace.

The AI adoption strategy shouldn't just be focused on cost and revenue but should also consider aligning AI investments with business and regulatory uncertainties. The rapid pace of emerging AI technologies has raised the question of whether investment in AI today would be obsolete or expensive a few months down the line, thereby hindering ROI clarity in AI projects. Indian banks are hence somewhat hesitant to ramp up adoption or investment in technological infrastructure.

Some of the other aspects include vendor lock-ins (considerable risk of lagging behind if the hyperscaler or vendor is slow to evolve). Lack of transparency in pricing and uncertainty of service consumption further add to investment concerns.

To mitigate these concerns, banks need to be strategic about how they design their AI systems and review their investment approach.

Designing for change: The ideal scenario would be designing modularised 'plug-and-generate' AI solutions where new LLM/SLMs can be updated with minimal interventions. AI solutions should be designed on open standards that make them portable and built with abstracted layers. These layers should be modularised for data/feature ingestions, AI model components (prompt interfaces, retrieval connectors, and evaluation frameworks), downstream workflows (app integration, feedback loops), model monitoring, and machine learning (ML)/ LLM operations (LLMOps).

Investing wisely: Compute power—AI chips/GPUs—is the backbone that drives AI. Investing in today's top-tier GPUs is capex-heavy. Enterprise GPUs degrade faster with increasing AI workload.⁹

Using GPUs effectively requires a deep understanding of how

they'll be applied. For inference workloads, it's relatively easy to estimate how many GPUs are required, whereas for training workloads the calculation is far more complex. A quick way to estimate the number of GPUs needed for inferencing would be to consider model size (parameters such as 7 billion or 70 billion), latency requirements (0.01 ms or 5 s depending on use-case), number of users, requests per-second during peak load. For training/fine-tuning, the variables to establish computing needs are not very well understood, and setting up such infrastructure is a highly costly affair.

For banks, committing to investing in fixed infra-assets would be an operational risk, both from a scalability and capability perspective. Over the years, we have seen increasing investments in emerging compute technology that is faster and cheaper and which directly competes with market giants.¹⁰

Our recommendations

1. Banks should opt for flexible, transferable computing strategies with GPU-as-a-service (GPUaaS) or rental models with security aspects as mandated in sectoral guidelines. This will ensure flexibility from vendor lock-in and new emerging tech perspective. Renting GPUs or using GPUaaS allows banks to scale up or down based on demand, avoiding idle hardware and reducing upfront capital expenditure.

Key takeaways:

- Build a separate repository for strategic and operational technologies.
- Since technology cycles are going to become shorter, organisations must plan for a shorter timeframe of three to five years for operational technology.

⁹ <https://www.cnbc.com/2025/11/14/ai-gpu-depreciation-coreweave-nvidia-michael-burry.html>

¹⁰ <https://www.ft.com/content/d8585870-17a5-43a0-95ef-cbebb1995107>

04


AI democratisation framework

Democratisation of AI would ensure that trust is built gradually. As AI is increasingly leveraged as an accelerator for routine tasks, technology will permeate the organisation and become a shared enabler rather than a privilege for select few, thus ushering in a transformational wave at an enterprise level.

Below is a proposed setup comprising three interconnected AI ecosystems that organisations can use to provide equitable access to this technology:

- **AI lab:** A federated environment with quick and easy access to cutting-edge technology, enabling rapid prototyping and experimentation which is accessible to a select group of AI specialists who drive innovation.
- **AI playground:** An approved and information-security complaint sandbox environment embedded in the current technology setup for building business critical AI solutions. This would be available to on-ground business units (BU) or function-tagged AI resources like data scientists, data engineers, DevOps, MLOps engineers, and Business analysts.
- **AI marketplace:** A platform that would host solutions developed in the playground. This user group would comprise business and operations teams who will extensively use AI features and AI-embedded products and solutions for downstream business operations and decision-making.

Proposed framework for democratisation of AI



These ecosystems would be overseen by the governance group—one with diverse skillsets which shoulders the responsibility of responsible AI adoption in the bank. This group would have end-to-end responsibility—from gatekeeping what kind of technology or solution is permissible

to ensuring the right kind of use cases are picked to minimise risk. The above framework provides a blueprint for fostering AI adoption across enterprises with guardrails on safe usage to ensure innovation thrives across all stakeholder groups.



05

AI adoption in India: Journey so far

India is building a sovereign AI ecosystem powered by government investment, indigenous innovation, and a banking sector which is ready to lead the transformation.

5.1 India’s AI momentum

\$524 million	\$17 billion	34,000+ GPUs	₹10,300 crore	#1 globally
GenAI startup funding in H1 2025 ¹¹	Projected AI market by 2027 ¹²	Sanctioned under IndiaAI Mission ¹³	Cabinet-approved IndiaAI Mission allocation ¹⁴	India’s rank in AI skill penetration ¹⁵

5.2 Indian AI platforms

Platform	Features
Sarvam.ai ¹⁶	Vernacular customer service, KYC automation, regional compliance
Bhashini ¹⁷	Multilingual AI support across 22 Indian languages for seamless customer engagement
Zoho AI (Zia) ¹⁸	Enterprise CRM and analytics with strong data privacy
Krutrim ¹⁹	Voice-based banking and multilingual chatbots
Meghdoot (C-DAC) ²⁰	Sovereign cloud for sensitive banking workloads

These indigenous solutions offer banks a path to AI adoption that aligns with data localisation requirements and DPDP compliance.

11

<https://enterpriseai.economictimes.indiatimes.com/news/industry/indian-genai-startups>

12

<https://www.business-standard.com/industry/news/ai-market-in-india-to-touch-17-billion>

13

<https://indiaai.gov.in/hub/indiaai-compute-capacity>

14

<https://pib.gov.in/PressReleasePage.aspx?PRID=2108810>

15

<https://www.pib.gov.in/PressReleasePage.aspx?PRID=2108810®=3&lang=2>

16

<https://www.sarvam.ai/>

17

<https://www.bhashini.ai/>

18

<https://www.zoho.com/zia/>

19

<https://www.kruti.ai/>

20

https://www.cdac.in/index.aspx?id=service_details&serviceId=MeghdootCloudSuite

5.3 India's integrated AI ecosystem

India's AI ecosystem operates as an integrated stack:


1. **Foundation:** Digital public infrastructure (Aadhaar, UPI) provides consented access to a trustworthy system.
2. **Datasets platform:** AIKosh provides seamless access to datasets (5000), AI models (250+) ²¹ and tools. It also provides access to a sandbox environment accessible through an **Integrated Development Environment (IDE)** (also recommended by RBI in its FREE-AI report).
3. **Compute:** IndiaAI Mission's GPU infrastructure and the Meghdoot app enable sovereign processing. ²²
4. **Models:** Indigenous language models (Sarvam, Krutrim, Bhashini) deliver India-first AI capabilities.
5. **Governance:** RBI FREE-AI and DPDP ensure safe and trusted AI adoption.
6. **Application:** MeitY has built a platform for developing and deploying AI solutions through the IndiaAI Application Development Initiative (IDAI).
7. **Research and innovation:** IndiaAI innovation centres are being launched across the country.
8. **Upskilling:** IndiaAI FutureSkills address the current gap between talent pool and AI skills.



21 <https://aikosh.indiaai.gov.in/home>

22 https://www.cdac.in/index.aspx?id=tap&tapcat=High_Performance_Computing

AI ecosystem template



Key insight: Layered sovereign infrastructure → Governed AI adoption → Banking outcomes at population scale

The bottom line: Indian banks have a unique opportunity to build on this infrastructure, leverage homegrown AI platforms, and operate within a governance framework designed for responsible innovation. The question that then arises is whether the banking ecosystem can figure out a way to ride this wave.

The building blocks are in place. Success will depend on how effectively banks are able to leverage government-backed initiatives for scaled impact.

06

Conclusion


Although AI has been around for a few decades, it started to be used widely by enterprises only recently, thanks to advancements in computing technology. Despite this rapid adoption of AI, several barriers need to be overcome in order to realise its full potential. AI has already begun reshaping the future of most industries, including the banking sector. Our research indicates considerable interest in this technology among the leadership and key stakeholders of India's banking industry. Several large private banks have integrated mature AI solutions in their key operational processes. Further, PSBs and other institutions are actively embracing their AI transformation journey through pilots or have included plans for the same in their immediate roadmap. However, concerns around reputational risks and loss of customer trust are key roadblocks to AI adoption at scale. Addressing these challenges head-on will be critical.

Banks will have to walk a very fine line between opportunity and caution. It is essential for the sector to be aware of the

various risks AI poses and set up robust governance and mitigation plans. Enterprise-wise commitment and strategic AI investment across evolving technology and use case selection should be embedded in banks' short- and long-term vision. This will be the pathway to ride the AI wave.

The Government of India has been driving AI adoption by both businesses and governments alike through initiatives such as the IndiaAI Mission, which focuses on holistic AI ecosystem development. The availability of AI infrastructure, transformative solutions, governance, data assets, and capabilities led and developed by the country presents a unique opportunity for Indian banks to accelerate their transformation.

By placing the right bets on scalable and robust technology, innovation, and governance, we believe the Indian banking industry can build trustworthy solutions, thereby aligning itself with the vision of realising Viksit Bharat by 2047.





About PwC

We help you build trust so you can boldly reinvent

At PwC, we help clients build trust and reinvent so they can turn complexity into competitive advantage. We're a tech-forward, people-empowered network with more than 364,000 people in 136 countries and 137 territories. Across audit and assurance, tax and legal, deals and consulting, we help clients build, accelerate, and sustain momentum. Find out more at www.pwc.com.

PwC refers to the PwC network and/or one or more of its member firms, each of which is a separate legal entity. Please see www.pwc.com/structure for further details.

© 2025 PwC. All rights reserved.

Contact us

Gayathri Parthasarathy

India Financial Services Leader and
Global Financial Services Technology Leader
PwC India
gai.partha@pwc.com

Vivek Belgavi

Partner and Leader, Financial Services Technology
PwC India
vivek.belgavi@pwc.com

Hetal Shah

Partner, Data and Analytics
PwC India
hetal.d.shah@pwc.com

Anit Shanker

Partner, Financial Services
PwC India
anit.shanker@pwc.com

Contributors:

Hetal Shah
Neeraj Sibal
Princia Vaz
Meenaaz Bastawala
Nikhil Joshi
Middepalle Somnath Reddy
Devdatta Biswas

Editorial:

Dion D'Souza
Rubina Malhotra

Design:

Harshpal Singh
Kirtika Saxena

pwc.in

Data Classification: DC0 (Public)

In this document, PwC refers to PricewaterhouseCoopers Private Limited (a limited liability company in India having Corporate Identity Number or CIN : U74140WB1983PTC036093), which is a member firm of PricewaterhouseCoopers International Limited (PwCIL), each member firm of which is a separate legal entity.

This document does not constitute professional advice. The information in this document has been obtained or derived from sources believed by PricewaterhouseCoopers Private Limited (PwCPL) to be reliable but PwCPL does not represent that this information is accurate or complete. Any opinions or estimates contained in this document represent the judgment of PwCPL at this time and are subject to change without notice. Readers of this publication are advised to seek their own professional advice before taking any course of action or decision, for which they are entirely responsible, based on the contents of this publication. PwCPL neither accepts or assumes any responsibility or liability to any reader of this publication in respect of the information contained within it or for any decisions readers may take or decide not to or fail to take.

© 2025 PricewaterhouseCoopers Private Limited. All rights reserved.

HS/December 2025 - M&C 50545