CXO SURVEY & REPORT 2026

CROSSING THE AI DIVIDE

What it takes to move from AI adoption to AI at scale in Indian banking

Contents

Foreword

AI adoption in Indian banking is advancing faster than the institutional readiness needed to support it.

Banks have moved beyond experimentation, with production deployments visible across a growing range of functions. Many leaders understandably view this as evidence that their institutions are progressing towards broader adoption. The capability-level evidence presents a more nuanced picture. The foundations needed to extend these successes consistently across the institution are developing at different speeds.

This is what we set out to examine. The study brings together perspectives from technology, data, operations and risk leaders across Indian banks and NBFCs. It looks beyond visible adoption to the data, infrastructure, workflows, engineering practices, governance mechanisms and workforce readiness beneath it.

The findings show meaningful progress. Banks have invested substantially in connecting their technology estates, improving access to data and deploying AI in selected functions. Governance, engineering controls and internal skills are less mature, making it harder to carry early successes into new functions without recreating much of the surrounding environment.

Wider adoption will also place new demands on core banking systems. Making services accessible through APIs was an important step. AI requires those systems to provide data and banking services with the context, permissions and responsiveness needed to participate safely in workflows and decisions. This must be supported by shared technology foundations and clear boundaries of authority, accountability, and oversight.

For banking leaders, the opportunity is to look beyond the growth in individual use cases and consider how readiness is developing across the institution. Understanding which capabilities require greater attention can help guide investment towards the foundations for the next stage of AI adoption.

We hope this report provides a useful framework for making that assessment and turning the progress already underway into sustained institutional capability.

Sivaram Kowta
President, Zeta

Preface

Artificial Intelligence has moved well beyond being a technology trend in financial services. It is becoming part of how banks serve customers, manage risk, run their operations, develop their people and make decisions. As banks move from experimentation to real-world adoption, the important question is no longer whether AI will change banking, but how deeply that change will go, how quickly it will happen, and what it will mean for the industry and its people.

This study looks at that transformation through the lens of the Indian banking industry. It explores the intersection of AI, financial technology, business processes and organisational capabilities, while looking beyond the current excitement around generative and agentic AI. The focus is on what is actually changing within banks: where AI is being adopted, which functions are being transformed, how roles and skills are evolving, and what capabilities banks will need as they move into the next phase of digital banking.

AI adoption is not simply a technology decision. Its success depends on how technology comes together with people, data, processes, governance, risk and business strategy. This study therefore looks at AI not only as a means of improving efficiency, but also as an opportunity to rethink how financial institutions work, serve customers and create value. The research brings together primary research, industry perspectives, executive inputs, secondary research and an examination of emerging developments in AI and financial technology. Rather than attempting to predict a single future for banking, the study looks for the patterns, opportunities and challenges that can help banking leaders understand the direction in which the industry is moving.

One of the central themes of this study is the importance of functional and domain knowledge. The future of AI in banking will require more than technical expertise. It will require people who understand banking products, processes, customers, risk and regulation, and who can combine that understanding with new technologies. The ability to bring deep financial services knowledge together with AI and technology could become an important source of competitive advantage.

We hope this study adds meaningfully to the conversation about AI in Indian banking, not simply by capturing where the industry stands today, but by encouraging leaders to think more deeply about the kind of bank they want to build for tomorrow and the capabilities needed to get there.

Manoj Agrawal
Group Editor, Banking Frontiers

Executive Summary

AI is live in Indian banking. The stack it runs on was not built to carry it.

Indian banks have moved AI into production. Customer service, fraud and risk analytics, document processing and software testing all report live deployments, and 70% of CDOs place their bank at selective or scaled deployment. The question for bank leadership is whether success in one function can be reproduced efficiently and safely in the next.

> Indian banks are deploying AI selectively, while the capabilities needed to repeat those deployments remain closer to pilot. The divide is not between banks that use AI and banks that don't. It is between piloting well and deploying repeatably.
>
> Core systems, processes and controls were designed before AI began participating in banking workflows. Much of the necessary context and control is therefore rebuilt for each deployment.

Self-assessed maturity is about a stage ahead of capability readiness

Seven in ten CDOs place their bank at selective or scaled deployment. We assessed each participating institution on five critical capabilities against the gate each one must clear. No capability is consistently at selective deployment. Governance, the weakest on the scale, is in an early pilot stage.

EXHIBIT 1 The confidence-capability gap

The two assessments are not contradictory. A CDO assessing the bank as a whole will reasonably weigh what is live. A gate-based assessment weighs what could be repeated. They are different questions, and right now they have different answers. The gap matters because AI investment decisions may reflect the breadth of adoption, while delivery depends on the readiness of the underlying capabilities.

Control readiness, not lack of conviction is constraining wider deployment

Taken one at a time, the headline numbers each support a reassuring story. However, taken together they point to three specific conclusions.

1. **Value is visible, but not yet consistently repeatable.** Seven in ten CDOs report AI live in at least some functions. Yet, four in ten could not identify a high-ROI use case in their own institution.
2. **Investment remains measured, but conviction is not the principal reason.** Most respondents put under 10% of new-project technology spend into AI, including banks already running deployments across several functions. Technology leaders nevertheless rate lack of ROI clarity as the lowest barrier, while executive scepticism and employee resistance rank below skills and security.
3. **Control readiness is the principal barrier.** Three-quarters of technology and risk leaders identify security and data privacy as the leading barrier. No respondent has an implemented Responsible AI framework. Only one CRO in five rates model-risk management approaching maturity, and no governance dimension reaches 3 of 5.

That constraint also explains the budget. Investment follows value that can be repeated, and repeatable value follows the ability to deploy safely without rebuilding the environment each time. Until successive use cases become less costly and easier to control, spending under 10% is the sensible answer.

The same two requirements recur across all five capabilities

We assess five capabilities against five stages of maturity. The table places each on the evidence and is separate from respondents' assessment of overall maturity.

Capability	Where it stands	The evidence	What it means for the stack
Data & infrastructure	Closest to selective deployment	73% have real-time data platforms and 73% API-first architecture; 40% have an MLOps platform. Data is “mostly ready” for 80%.	Access is stronger than AI usability. Context and permissions do not yet travel consistently with the data.
Operations & workflows	Approaching selective deployment	88% report impact in retail lending, 75% in customer service. Four in ten cannot name a high-ROI use case.	Impact sits at the task, not the workflow. The second deployment can't reuse the first.
Engineering & SDLC	Mid-pilot	80% use AI for test generation; 30% for deployment. Security and privacy rate far above lack of ROI clarity as a barrier (3.89 vs 2.00 of 5).	AI can write, not act. The lifecycle has no safe way to let it change anything.
Governance & trust	Early pilot	No implemented Responsible AI framework; 20% of CROs rate model-risk management very mature; every dimension below 3 of 5.	Policy development is ahead of enforcement. AI identity, permissions and audit mechanisms are still developing.
Workforce & organisation	Mid-pilot	Talent rates as the top contributor to outcomes and internal reskilling as the last, at the floor of the scale (3.88 vs 1.00 of 5).	External capacity is advancing faster than internal reskilling. Domain-team fluency remains limited.

Two implications follow. The weakest capability sets the pace, and governance and engineering are exactly what the move from selective to scaled depends on. More use cases will not open that gate. Even the strongest capability is short of it. Connectivity was the priority of the digital decade whereas repeatable AI also requires context, permissions and a consistent operating environment.

API adoption alone (at a high 73% in the study) should therefore not be read as AI readiness, because the gaps become more visible as deployments multiply.

Across the five chapters the same two foundational requirements recur.

- **Banking systems AI can use with context.** An AI agent can now access a limits API. However, it cannot necessarily tell what that limit means, how it relates to a delinquency bucket, or under which policy it may act on either. That meaning has to travel with the data, along with the consent and permission rules that govern its use.
- **A shared control plane for AI.** As AI moves from generating outputs to calling services or changing systems, the institution must know which model or agent is acting, what it may access, which policy applies, where a human approval is required and what has been recorded. Operations, engineering and governance each require the same shared layer.

Neither is a use case. Both determine whether successive deployments become faster and less expensive.

Five priorities for repeatable deployment

Each addresses a capability gate.

1. **Measure repeatability.** Instead of just the number of live use cases, measure by the cost and elapsed time of successive deployments, alongside the number of live use cases. Repeatability provides stronger evidence of maturity than breadth alone.
2. **Make core systems usable by AI, not merely accessible.** Extend API access with real-time, event-driven data carrying the context and permissions AI needs to interpret limits, delinquencies and other banking concepts correctly.
3. **Establish the control plane before consequential deployments multiply.** Agent identity and permissions, policy enforcement, defined approval points, and auditability convert Responsible AI principles into enforceable controls and provide a basis for risk and security approval.
4. **Build shared platform components.** Data products, model lifecycle, engineering context, and integration patterns should be built once and reused. An initial deployment may be bespoke; subsequent deployments should draw increasingly on shared foundations.
5. **Fund the shared foundations.** The sub-10% allocation moves on its own once the marginal deployment is cheap and controllable. Priorities should be sequenced according to the least mature capability.

The AI Maturity Framework

Five stages, five capabilities, and the rule that the weakest one sets the pace

AI adoption in a bank does not follow one curve. You can run a fraud model in production while your Responsible AI framework sits in draft. You can hire specialists faster than you reskill the teams around them. Any honest assessment of AI readiness has to hold those differences apart rather than average them into a single score.

Each stage is a stronger claim than the last

Stage	Where it stands	What it looks like
Exploratory	Proofs of concept and sandboxes, with little or no production exposure.	-
Pilot	AI live in bounded use cases, built one at a time and closely watched.	Prove relevance. A bounded use case shows credible fit and access to the data it needs.
Selective deployment	Several use cases running in chosen functions under established controls but repeating them is slow.	Prove value and control. It works in production under defined data, security, governance and oversight controls.
Scaled deployment	Deployments repeat on shared data, engineering, operating and governance capabilities.	Make it repeatable. The bank can reproduce a deployment without rebuilding the environment around it.
Institutionalised	AI is part of how the bank runs, introduced through established platforms, controls and roles.	Embed the capability. Roles, governance, engineering and workflow design support it as a matter of course.

Across these stages, the leadership question moves from what value AI can create to can it run under control, then can we repeat it, and finally is this how the bank now works. Most of the industry is standing at the third question, one of repeatability.

Readiness must be assessed across five capabilities

1. **Data & infrastructure -** whether AI can reach the data and banking capabilities it needs, and whether the technology estate can run it reliably.
2. **Operations & workflows -** whether AI can move past assisting single tasks and take part in how the work gets done.
3. **Engineering & SDLC -** whether AI can be used across how technology is designed, built, tested, deployed and run.
4. **Governance & trust -** what AI may access, recommend, decide or execute, and under whose authority.
5. **Workforce & organisation -** whether people can direct, supervise and work alongside increasingly capable systems.

Leadership conviction and the partner ecosystem help across all five, and neither showed up as a binding constraint in our study. Progress slows when one capability falls behind what the next stage needs. For instance, strong data availability does not compensate for uneven readiness in governance.

No capability has cleared the gate to scaled deployment yet

Two exhibits carry the placement of each capability, and we use them consistently for the rest of the report. The matrix shows which stage each capability has reached and what that stage looks like in practice. The gate chart shows where each sits inside that stage, and which gate it faces.

EXHIBIT 2 The AI capability maturity matrix

	EXPLORATORY	PILOT	SELECTIVE DEPLOYMENT	SCALED DEPLOYMENT	INSTITUTIONALISED
Data & infrastructure	Data pockets; sandbox compute	Use-case datasets; cloud experiments	Governed domains; real-time platforms and APIs • TODAY Just entered	Reusable data products; hybrid AI infrastructure	Governed, contextual data available systematically for AI
Operations & workflows	Individual productivity and task experiments	AI assists bounded operational tasks • TODAY At the gate to selective	AI contributes across selected workflows; hand-offs explicit	End-to-end workflows redesigned around human and AI	AI systematically embedded across operating processes
Engineering & SDLC	Individual developer tools	AI used in bounded engineering tasks • TODAY Mid-stage	AI adopted across multiple SDLC stages under defined controls	Integrated across build, test, deploy and operate	AI embedded across the lifecycle with systemic controls
Governance & trust	Policy discussions; frameworks nascent	Case-by-case approvals; Responsible AI framework drafted • TODAY Just entered	Model-risk practice; defined ownership	Decision-lifecycle controls; AI identity and permissions	Controls embedded by design; autonomy expands within boundaries
Workforce & organisation	Enthusiast-led experiments	Specialist hires; external partners • TODAY Mid-stage	Centres of excellence and cross-functional pods	AI-enabled roles; systemic reskilling	AI fluency embedded across relevant roles

Marked cells show where the capability evidence places the industry today. They are not respondents' self-assessment of overall maturity.

EXHIBIT 3   Where the five capabilities stand against the gates

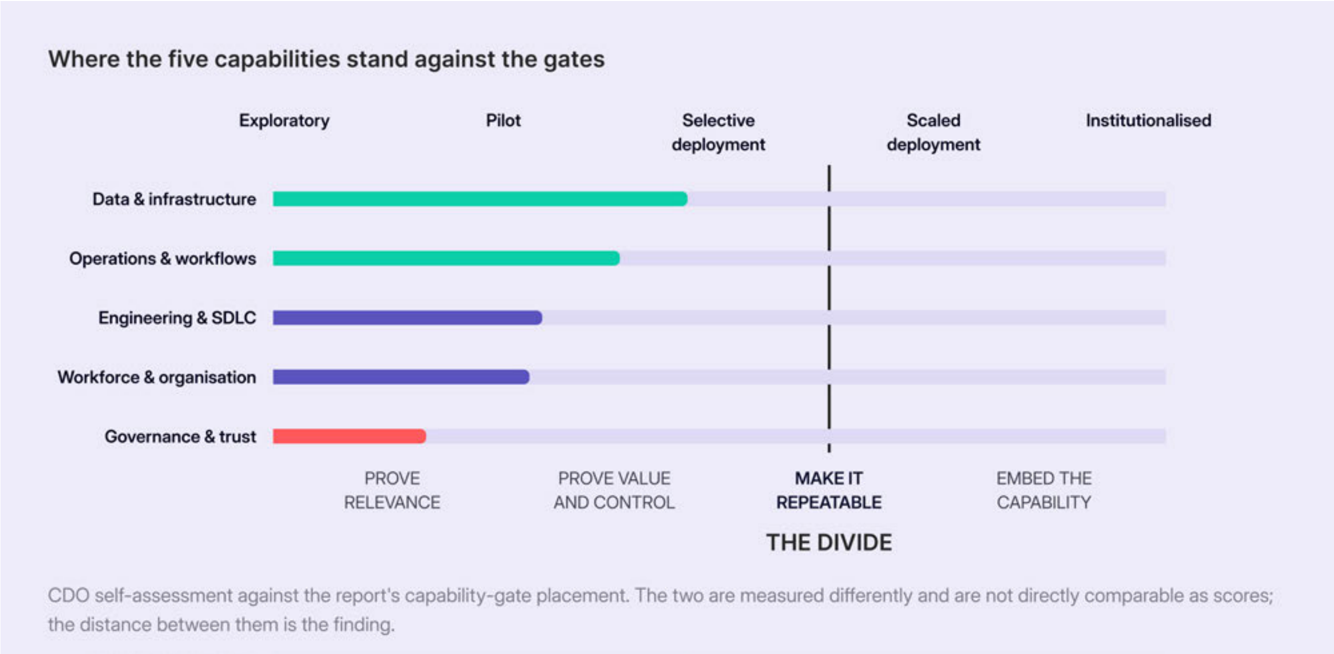

Sorted by maturity. The five chapters that follow run in stack order, working up from data; in gate order, governance comes first. Each opens with where the capability stands, sets out the evidence, and closes with what to build once instead of per use case.

A pilot is bounded and sits inside controls you already have, so an uneven stack is survivable. Scaled deployment is not, and at that point the weakest capabilities set the limit. Adding use cases does not move them: a use case does not necessarily elevate governance maturity, or give the engineering lifecycle a safe way to let AI act.

> **The divide is a cycle. One-off deployment keeps marginal cost high; high marginal cost keeps value local; local value keeps the budget measured; a measured budget leaves the shared layer unfunded.**

The matrix tells you where the one-off effort is going.

HOW TO READ THIS REPORT

The survey numbers are used to interpret readiness. The numbers cited in this study relate to adoption by function, maturity self-ratings, barriers, priorities and spend. The maturity framework uses this evidence to assess how far each institutional capability has progressed.

The interpretation is evidence-led. Respondents were not asked specifically about semantic layers, agent identity or policy enforcement. These requirements are inferred from the response pattern and our work with banks and point to what must mature for AI to move from selective deployment to repeatable scale.

Bases vary and several are small. Percentages should be read as directional. Floor and ceiling effects are noted where relevant.

Data & Infrastructure

Data is ready and trusted. What it lacks is meaning, permission and speed

TODAY:

Just into selective deployment. The most advanced of the five capabilities, but reachability does not yet mean AI usability.

The banking core has been connected. That was the work of the digital decade, and it is largely done. It made banking capabilities reachable, but not yet fully usable by AI systems that need context and permission to participate in a workflow.

An API tells an agent a limit exists. It does not explain what the limit means, how it relates to a delinquency bucket, or under which policy the agent may act on either. Without that context, each deployment has to rebuild meaning and controls around that use case. This is the pattern in the constraints CIOs identify.

Data usability ranks highest among constraints; availability and quality are not concerns

The instinct is to treat AI readiness as a data-quality programme. Whereas the problem sits elsewhere. Eight in ten CIOs and CTOs call their data “mostly ready” for AI at scale, and the problems they name are about usability, not accuracy.

EXHIBIT 4 The data gap is about usability

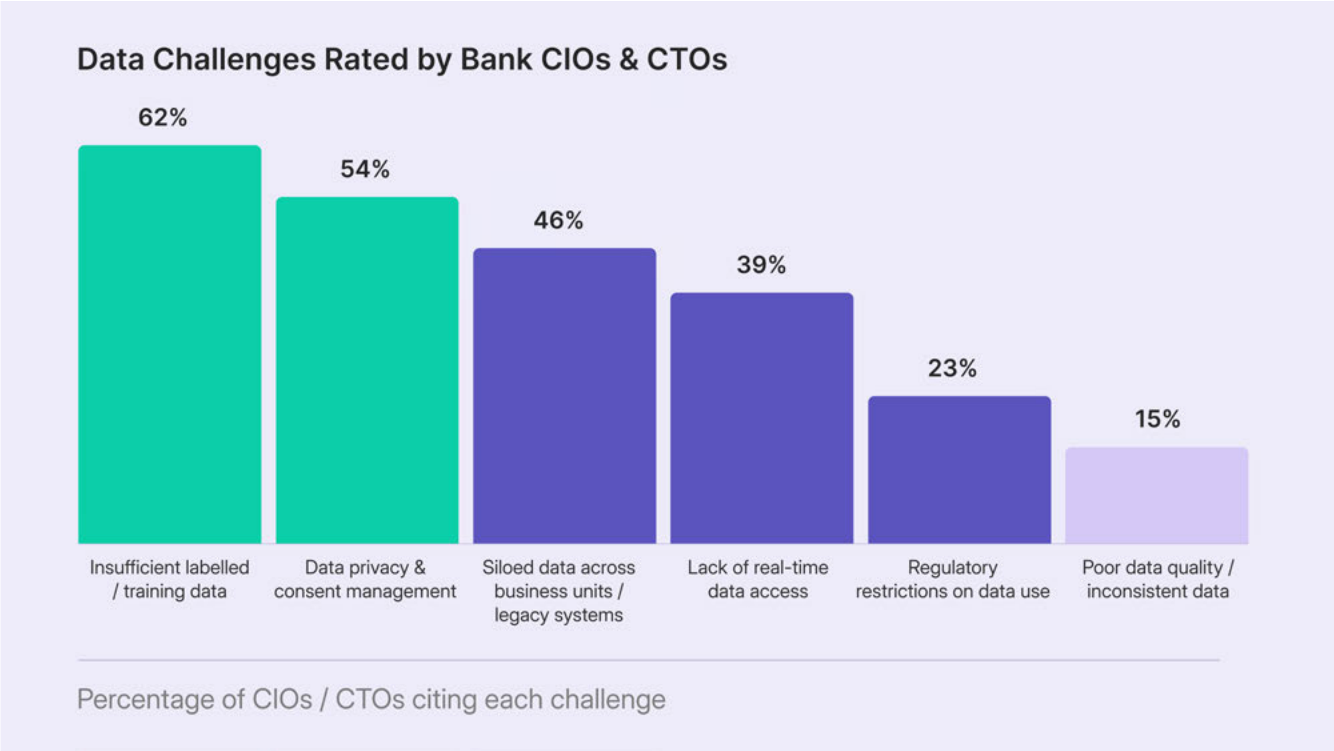

62% cite insufficient labelled or training data and 54% cite privacy and consent management. Siloed data follows at 46% and lack of real-time access at 39%. Poor data quality comes last, at 15%.

The pattern is important. The data exists, sits in silos, and is broadly trusted. What it lacks is meaning an agent can act upon, permissions it can check, and access at the speed the decision needs. Those are design gaps, and today they get closed by hand, once per use case.

Banks are already pointing AI at the problem: 56% use it to enrich data and a further 11% have a proof of concept running, covering segmentation, spend analytics, training-data discovery and data dictionaries. The risk is that it stays inside individual projects, so enrichment done for one use case gets done again for the next.

Four data attributes belong in the shared layer, not in each deployment

The move is from managing data as an asset to serving it as a product - governed, reusable, and usable by many applications at once.

1. **Context & metadata -** what the data means in banking terms, where it came from, how it relates to everything else.

2. **Enrichment & labelling -** the classification and derived attributes that make raw data useful to a model.
3. **Permissions & consent -** purpose, consent and access rules carried with the data, checkable at the point of access.
4. **Quality, lineage & access -** trustworthy, traceable, and available at the freshness the use case needs.

Take next-best-action for an existing customer as an example.

Here is what changes when these four attributes are in place.

Data need	Today	With a shared data layer
Customer context	Data sits across product and channel systems	Signals assemble into one current, governed view
Enrichment	Raw data carries little meaning outside its source	AI classifies, derives attributes and adds context
Permission	Consent and access checked application by application	Purpose, consent and access rules travel with the data
Freshness	Batch extracts or a pipeline built for this use case	The data product exposes the view the decision needs

Connectivity is ahead of AI operating readiness

EXHIBIT 5 Connective foundations lead the AI operating layer

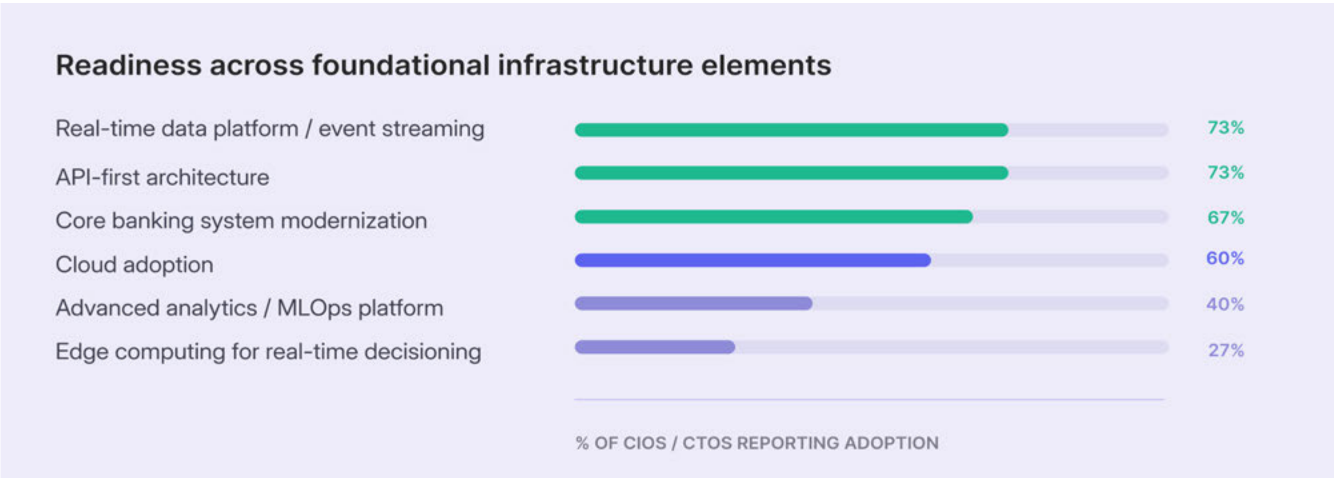

CIOs and CTOs report strongest progress in the connective foundations: 73% cite real-time data platforms, 73% cite API-first architecture, 67% cite core banking modernisation and 60% cite cloud adoption. The operating layer for repeatable AI is less mature: advanced analytics and MLOps stand at 40%, while edge computing for real-time decisioning stands at 27%.

The first four numbers suggest that much of the banking estate is now reachable. The fifth shows that the operating layer for repeatable AI is still less mature. Until banks have a consistent way to run AI workloads on top of this connectivity, each use case remains more bespoke than reusable, and the marginal deployment stays expensive.

AI is also a different kind of workload. It needs compute matched to latency and data-residency rules, models released through one lifecycle, security that follows the model rather than the deployment, and monitoring that watches model behaviour and cost as well as uptime. Most teams handle each of those per deployment. Each belongs in the platform every AI workload inherits.

- **Compute & workload placement -** each workload in the right approved environment, at the right capacity and latency.
- **Model deployment & lifecycle -** one process to deploy, version, update and retire.
- **Security & isolation -** protection matched to the sensitivity of each workload.
- **Performance, resilience & cost -** availability, latency, usage and cost watched together.

THE GATE AHEAD

AI can find and use enterprise data with the banking semantics, permissions, lineage and freshness it needs, without rebuilding any of that for the next application. AI workloads scale and change across approved environments through a common lifecycle, security model and observability layer.

Operations & Workflows

AI is helping with tasks. The next step is redesigning the workflow around it.

> **TODAY:**
>
> At the gate between pilot and selective deployment. AI is delivering real impact in bounded operations, added to individual tasks inside processes built around human hand-offs, with each deployment integrated and controlled on its own.

Banking operations run on systems designed for people: screens, queues, work items, approvals passed from one desk to the next. AI has entered at the task level, where it reads a document, drafts a response or makes a recommendation and hands the result back to a human.

That is where the impact is today, and it is also the ceiling. To move from assisting a task to progressing a case, core systems have to expose banking capabilities as services rather than screens, events have to trigger work rather than queues, and the workflow has to say what AI may do and where a person decides.

AI impact is strongest where work remains reviewable

88% of COOs report meaningful impact in retail lending and 75% in customer service. CASA and back-office operations follow at 63% each, fraud operations at 50% and trade finance at 38%.

EXHIBIT 6 *Where AI is creating operational impact today*

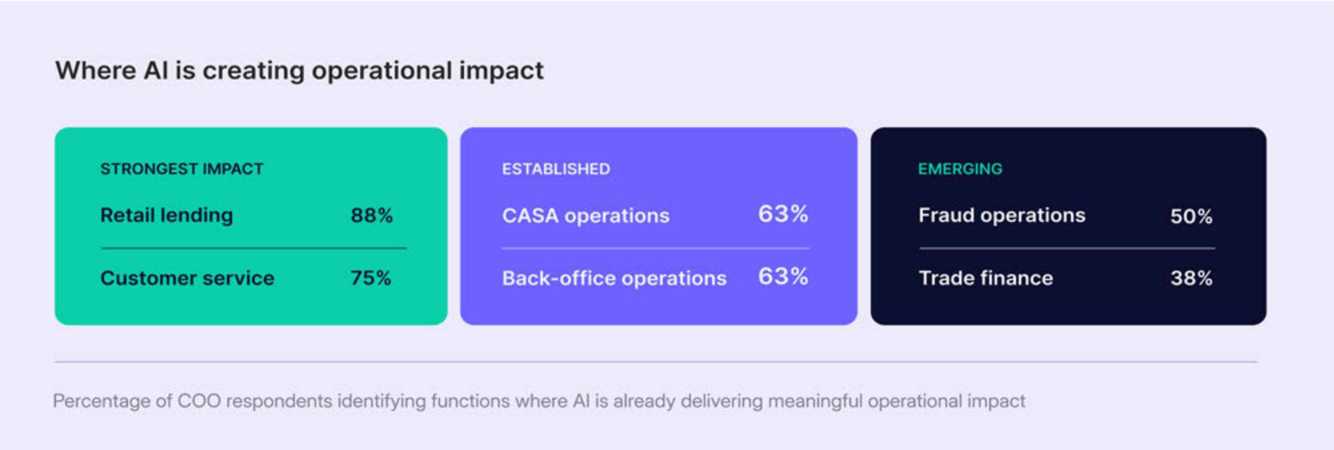

AI impact is showing up first in workflows with high volume, repeatable processing and clear human review. These are good conditions for pilots, but they do not by themselves prove that the same impact can be reproduced across workflows. Retail lending and customer service have more of these conditions; trade finance, with unstructured documents and case-by-case judgement, has fewer.

The tension is visible in the survey. COOs report strong impact in retail lending, yet roughly four in ten CDOs could not identify a high-ROI use case in their own institution. The implication is that impact is being felt at the task level, while value is not yet being established consistently at the workflow level.

Redesigning the workflow means redesigning what the core exposes

An agent progressing a credit-card application has to call eligibility, KYC and bureau checks as services, receive the events that say a document arrived or a limit changed, understand what those results mean in your bank's terms, and act only inside the policy you set. Where the core offers screens and batch files instead, every workflow integration gets built by hand.

Expose capabilities that way and AI can take part across many workflows without a new integration each time. The controls around it get inherited rather than rebuilt.

Take credit-card origination as an example.

Routine, reversible steps need light oversight. Exceptions and consequential calls need more.

Origination task	Today	With AI participating
Application intake	Capture, validate, follow up	Assemble context, spot gaps, progress complete cases
Data & checks	Separate KYC, bureau and eligibility steps	Coordinate the checks and consolidate outcomes
Credit decision	Rules, models and analyst review	Combine signals and act inside policy
Human review	Spread across multiple stages	Focused on exceptions and overrides
Progression	Sequential hand-offs	Continuous progression with exception routing

Controls should be designed upfront before deployments multiply

Once AI starts calling services, progressing cases or executing decisions, the control model must be designed upfront. Interfaces, permissions, policy controls and traceability need to be reusable across workflows, not rebuilt for each deployment.

EXHIBIT 7 What safe, repeatable AI participation needs

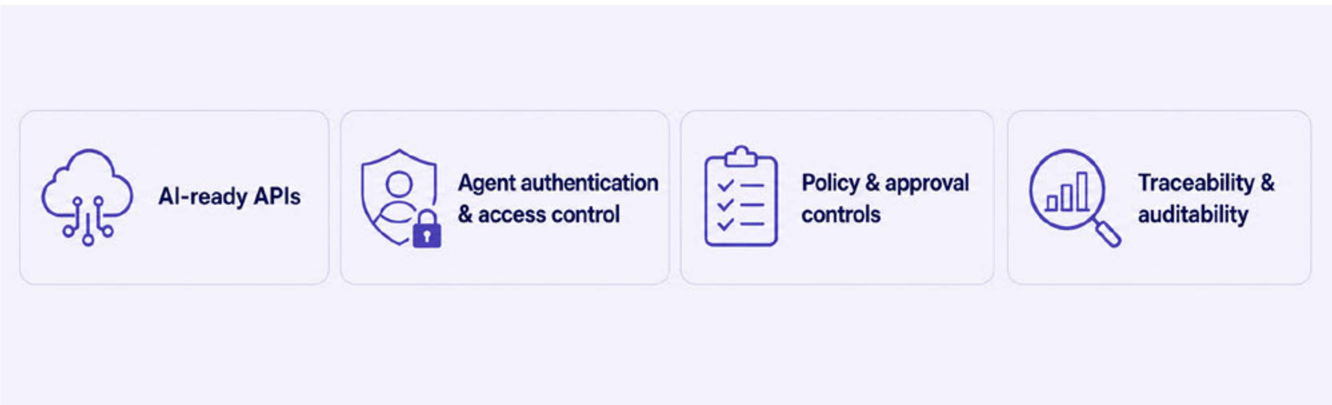

- **AI-ready APIs -** controlled access to banking capabilities instead of point-to-point integrations.
- **Agent authentication & access control -** which machine is acting, and what it may reach or run.
- **Policy & approval controls -** actions held inside business rules, thresholds and required signoffs.
- **Traceability & auditability -** the context, recommendation, action, policy applied and any human intervention, preserved.

The last three are the same controls engineering and governance need. Build them for operations and you have most of what the other two require.

THE GATE AHEAD

AI shapes how work is coordinated and how decisions get made, on a core that exposes capabilities and events, with controls AI inherits rather than rebuilds.

Engineering & SDLC

Strong where AI produces reviewable outputs. Thin where AI would alter code, releases or production systems

> **TODAY:**
>
> **Mid-pilot.** AI is widely used in testing and code generation. Adoption falls away sharply wherever AI would change or execute rather than produce something a person reviews.

Test generation and QA lead at 80% of CIO and CTO respondents, then code generation at 60% and code review at 40%. Specifications and documentation, deployment and CI/CD, and incident detection and root-cause analysis each sit at 30%.

EXHIBIT 8 AI penetration across the SDLC

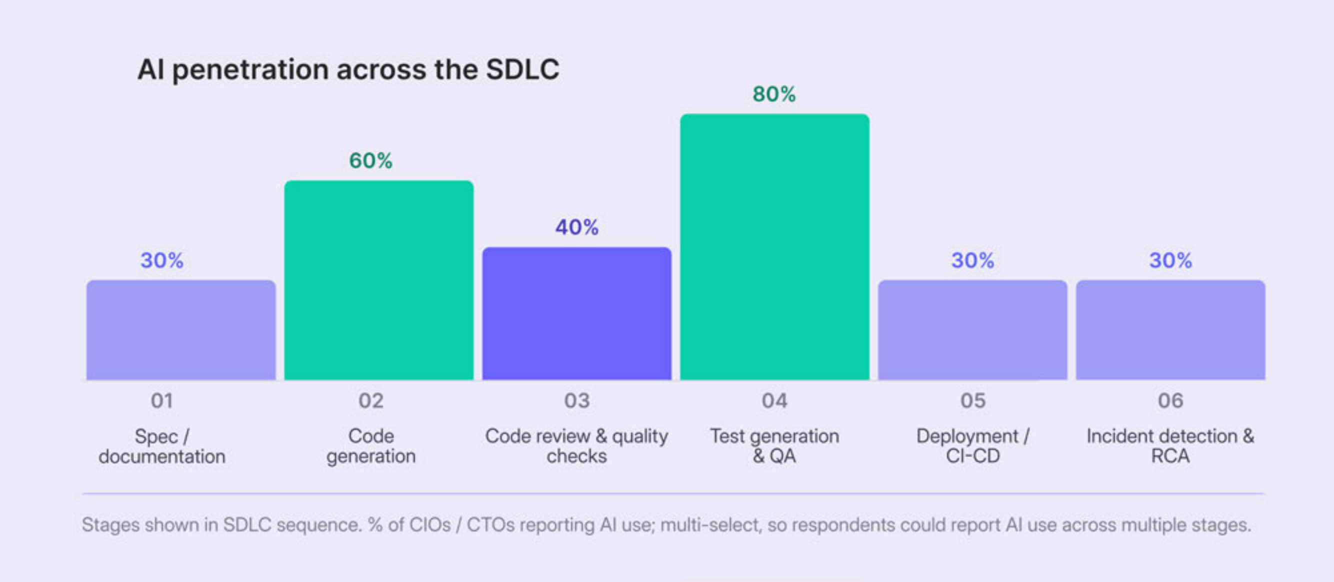

The pattern is not about difficulty but about consequence. The high-adoption stages produce something a person reads before anything happens. The low-adoption stages change a repository, trigger a pipeline or act on production.

Engineering leaders are not unconvinced; controls are still developing

Technology leaders rate security and data privacy far above lack of ROI clarity as a barrier (3.89 vs 2.00 of 5). Conviction is not the constraint. The absence of a way to bound what AI may change is.

So the question shifts from whether AI can help an engineer to what AI can see, change or run - and under what approval, security and accountability.

> *"As AI accelerates software creation, the strategic IP extends beyond generated code to the business knowledge, product design, architecture, data models, security controls and integration framework that define how the bank operates."*
>
> **Dr. Rajendran N.**
> *Former CTO, NPCI and Former CEO, IFTAS*

The gain comes from context carrying across the lifecycle

AI is already making code generation, review, testing and documentation faster. The larger gain arrives when requirements and architecture context carry into review and testing, and a production signal traces back to the change that caused it.

Take a regulatory change as an example.

The shift is continuity of context, not autonomous delivery. Human and system controls stay where they are.

Stage	Today	With AI across the lifecycle
Requirement & impact	Teams interpret the change and work out which applications are affected	AI assembles the context and maps likely code, service and dependency impact
Build & review	Engineers implement; review happens separately	AI proposes changes and reviews them against coding, security and architecture standards
Test	Teams write and run regression tests around the change	AI generates tests from the requirement and the components it touches
Release	Approved changes move through CI/CD and release gates	AI prepares and progresses release steps inside existing approvals
Operate	Teams diagnose from logs, alerts and service context	AI correlates signals, proposes root cause and recommends a fix

SDLC needs a safe way to let AI act

AI use across the lifecycle remains fragmented if every team separately arranges code access, sets controls, defines approvals and creates an audit trail. These controls belong in the shared engineering environment, alongside the control plane described in operations and governance.

- **Engineering context & access -** the requirements, code, architecture and tools relevant to the task, under control.
- **Secure software controls -** existing security, quality and dependency standards applied to AI-written or AI-modified code.
- **Approval & execution controls -** what AI may change or run alone, and where a person must sign off.
- **Traceability & monitoring -** what AI generated or ran, connected to what happened downstream.

THE GATE AHEAD

AI is no longer concentrated in individual tasks and can work across multiple SDLC stages under established controls - which is the same investment governance is waiting for.

Governance & Trust

Policy is not the gap. Enforcement is.

TODAY:

Just into pilot, and the least mature of the five capabilities. Banks know they need AI governance and have policy in draft. The mechanisms that would enforce it inside the stack largely do not exist.

Governance is being built on two fronts, and both are early. On the risk side, banks are still working out how to validate models that learn, explain a decision an agent made, fix accountability and watch behaviour after release. Detailed supervisory guidance is awaited, and banks are right not to hard wire what may change.

On the systems side, a framework needs somewhere to run. The banking environment must know which system is acting, what it may reach, under which policy, and what was recorded. Most bank environments answer that for people and applications. Very few answer it for models and agents.

Governance has not yet reached operating maturity

EXHIBIT 9   AI governance capability - current maturity

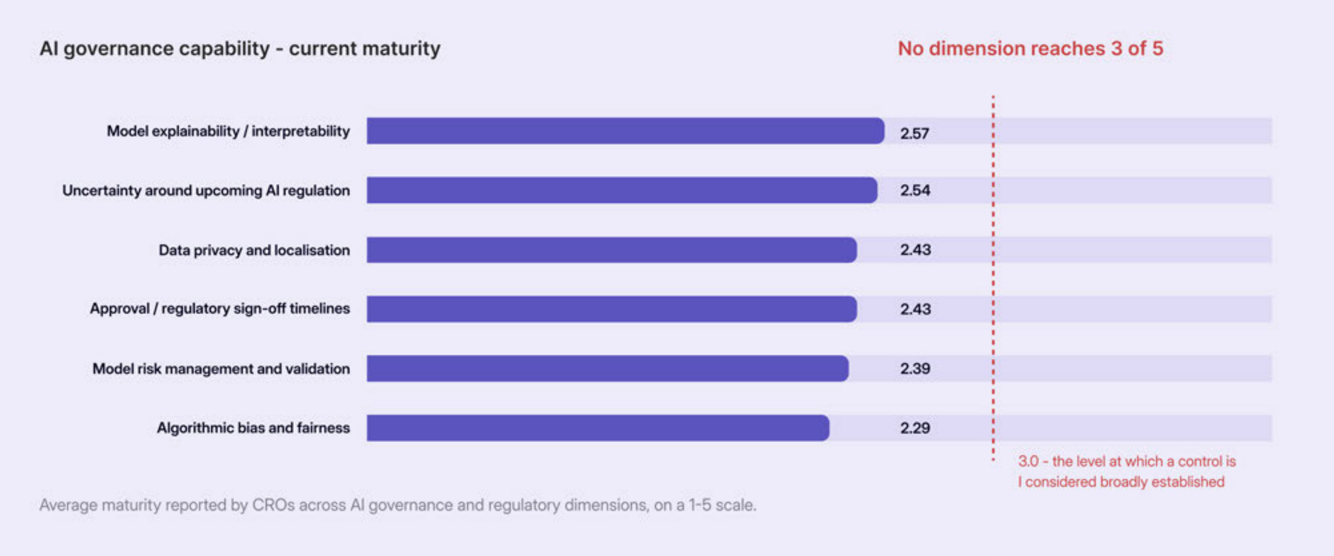

Model explainability rates highest and algorithmic bias lowest, with less than three-tenths of a point between them (2.57 vs 2.29 out of 5). This is not one weak area pulling down an otherwise mature control environment. It points to governance maturity that is still developing across the board.

The institutional picture matches. Only 20% of CROs call model-risk management approaching maturity. Responsible AI frameworks are in development at 60% of institutions and not yet on the roadmap at the other 40%, which is why nobody reports one implemented. Meanwhile at least 60% of CROs rank AI-led credit-risk models, early-warning systems and real-time fraud decisioning as top priorities for the next 18 to 24 months. Ambition is ahead of both the practice and the enforcement layer, and the gap is widening.

Security and privacy controls are the prerequisite, not the obstacle

Three-quarters of technology and risk leaders name security and data privacy as their leading barrier, and technology leaders rate it far above lack of ROI clarity (3.89 vs 2.00 of 5).

That is a control gap, not a lack of conviction. Banks are being asked to approve systems that may use sensitive data, influence decisions and act inside operational workflows before the surrounding controls are mature. Case-by-case approval is the responsible response until those controls exist.

Two things make that caution reasonable. Regulatory expectations are still moving, including the Reserve Bank's FREE-AI framework and future guidance on accountability, explainability and human oversight. And banks are implementing the Digital Personal Data Protection Act (DPDPA) while they scale AI.

Treat DPDPA and AI as one control build, not two. Consent, purpose limitation, retention and data-principal rights have to be recorded with the data and checked at the point of use. Those are the same controls AI needs when a model or agent assembles customer context, generates an output or acts inside a workflow. Building them twice is the expensive path.

Governing the model is no longer enough once AI acts

Traditional model governance asks whether a model is approved, validated and performing as expected. As AI moves closer to decisions and actions, governance also has to say what it may reach, recommend or run, and who stays accountable. Security controls must extend from users and applications to the AI systems themselves - identity, task-based permissions, policy enforced at the point of action, protection against manipulated inputs and leakage, and a trace of everything.

Take AI-assisted credit underwriting as an example.

Four things have to be governed, and the model is one of them.

What must be governed	Today	With AI participating
Model	Models validated and monitored on their own	Models and AI components governed through one lifecycle
Decision context	Rules define which data and factors enter underwriting	AI can use broader context inside defined data and policy boundaries
Authority	Credit policy defines human and system decision rights	AI recommends or decides only inside explicitly assigned authority
Outcome	Decisions and overrides recorded for review	Inputs, reasoning, actions and interventions all traceable

Model controls, permissions, approvals and audit belong in one layer

Building these separately for every use case gets harder to sustain as adoption grows. This is the control plane from the framework chapter, seen from the risk function: the same identity, permission, policy and traceability layer operations and engineering need, with model governance on top.

- **Model governance & validation -** models, data, performance and change across the AI lifecycle.
- **AI identity & permissions -** what each AI system is, what it may reach, what it is authorised to do.
- **Decision & policy controls -** business rules, risk limits, approvals and human escalation applied to AI decisions.
- **Traceability & oversight -** inputs, outputs, decisions, actions and interventions recorded so outcomes stay explainable.

Built alongside DPDPA rather than after it, these are what let a CRO sign.

THE GATE AHEAD

Controls extend from individual models to the decisions and actions AI takes part in, with clear ownership, permissions, policy boundaries and traceability. Until then a Responsible AI framework has nothing to enforce itself in, and security and privacy will keep being named as the leading barrier - because that is the honest answer.

Workforce & Organisation

Capability is being bought faster than it is being built in-house

TODAY:

Mid-pilot. Banks are adding specialists and starting to reshape roles and teams. Systematic reskilling of the existing workforce has barely started.

AI is changing work before it changes workforce size

Half of operations leaders expect AI-led productivity gains to free up capacity for higher-value work. 17% expect a broadly neutral effect and another 17% expect AI-enabled expansion. Nobody expects headcount reductions above 20%.

EXHIBIT 10 The near-term workforce effect is redeployment, not reduction

The subjective responses from COOs during this study point the same way - less manual effort, better productivity, more capacity - rather than roles disappearing.

Skills are the top constraint, and banks are relying on external capacity

CDOs rate talent availability and skill development as the single biggest contributor to AI uplift over the next two years (3.88 of 5). Operations leaders put skills shortages above employee resistance and conservative leadership (3.50 of 5). Capability constrains more than conviction does.

EXHIBIT 11 Banks are filling capability gaps externally faster than they are reskilling internally

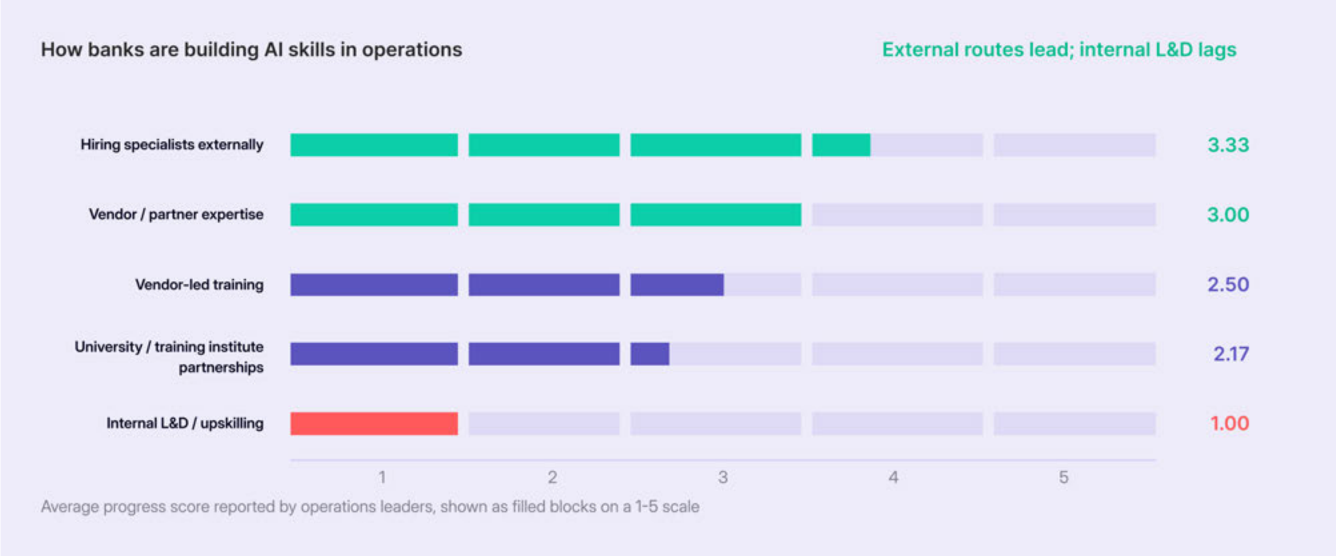

Operations leaders report considerably more progress through external hiring and vendor expertise than through internal learning and upskilling (3.33 and 3.00, compared with 1.00 out of 5). The internal L&D score sits at the floor of the scale, with every respondent selecting the lowest available rating. Given the small base, the precise score should be read cautiously; the more significant finding is the consistency of the response.

External expertise cannot replace institutional knowledge

Specialist teams can provide early capability. Broader fluency comes from the domain teams taking part - framing the tasks, judging the outputs, holding the call on decisions that matter.

Specialist hires and partners can build the platform. They cannot define how an institution interprets a limit, handles an exception or permits an override. That knowledge resides with the people who own its products, processes and risk decisions. Reskilling the domain workforce is therefore part of the technology programme, because it provides the context within which AI must operate.

Structures are starting to move too: pod-based teams replacing vertical ones, frontline engineering teams reorganised, AI centres of excellence spanning product, risk, data, technology and operations. Directional evidence rather than broad, but it is there.

"Sustainable AI adoption at scale requires technology professionals who combine banking knowledge, architecture, systems, data, security and operations with AI capability and strong governance."

Balaji Rajagopalan

CTO, State Bank of India

THE GATE AHEAD

AI capability reaches past specialist teams and external support into the roles and structures of the wider organisation - far enough that domain teams can specify what AI may mean and do, not only supervise what it produces.

The Divide by Institution Type

Broader adoption does not always coincide with stronger technology foundations

No institution type leads on every measure. Large private banks report the broadest production adoption, while mid-sized private banks report stronger technology-foundation maturity. Public-sector banks occupy the middle ground, and NBFCs remain earlier on both measures. As the respondent bases become small at this level of segmentation, these results should be treated as directional rather than as a ranking of institution types.

EXHIBIT 12 Breadth of production adoption against technology-foundation maturity

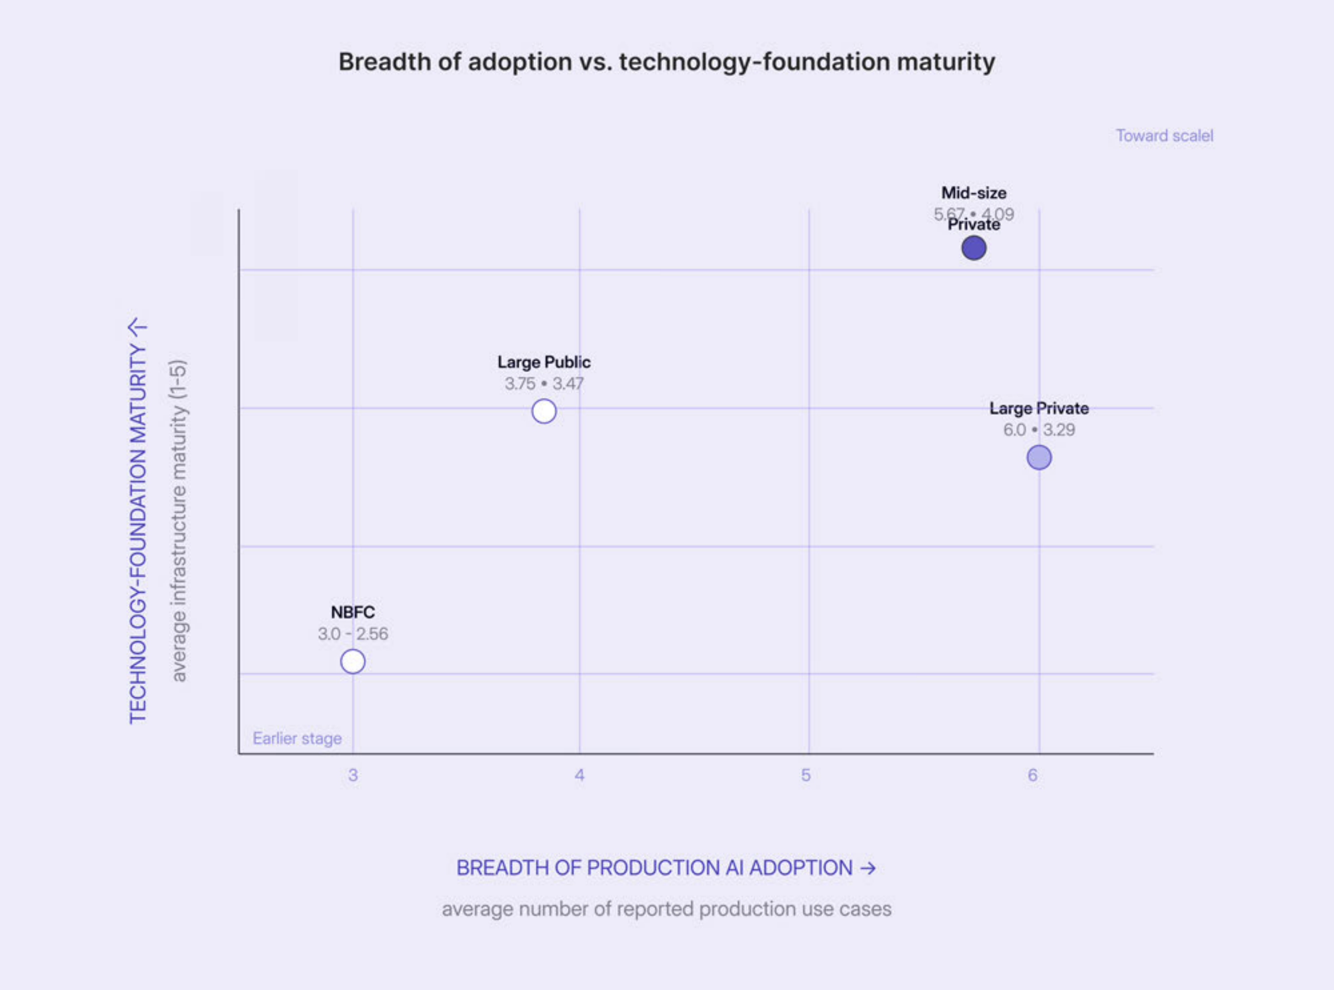

Large private banks report the broadest production footprint, but the lowest technology-foundation maturity among private banks: 6.0 use cases against 3.29 out of 5. **Mid-sized private banks** report nearly as many production use cases, supported by stronger foundations: 5.67 against 4.09.

This is the central segment-level finding: broader production adoption does not necessarily coincide with stronger foundations. For large private banks, the priority is improving consistency across deployments. For mid-sized private banks, it is translating stronger foundations into broader adoption.

Large public-sector banks show the reverse pattern, with stronger foundations relative to a narrower production footprint: 3.75 use cases against 3.47. Their opportunity is to replicate existing deployments more broadly. **NBFCs** remain earlier on both measures, at 3.0 use cases and 2.56 for foundation maturity.

> The useful question is not which institution type is most mature. It is which capability will constrain its next stage - and for the segment with the most AI in production, that capability is not the one it has been investing in.

Conclusion

Adoption is running ahead of institutional readiness

Indian banks are no longer at the beginning of AI adoption. Production deployments are visible across customer service, fraud and risk, lending, operations and software engineering. Seven in ten CDOs place their banks at selective or scaled deployment.

Assessed capability by capability, the picture changes. Data and infrastructure have crossed into selective deployment. Operations sit at the gate to it. Engineering and workforce are mid-pilot. Governance has only just cleared the first gate. None of the five is at the gate that matters, and the two furthest back are the two that gate the crossing.

> **AI being live is not the same as being ready to scale.** Banks have shown it works in production. The next test is whether those successes can be repeated across the institution without rebuilding the supporting capabilities each time.
>
> **That is the divide.**

Crossing the divide is not primarily a question of appetite. Lack of ROI clarity is the lowest-rated barrier, while executive scepticism and employee resistance rank below skills and security. The evidence points instead to a structural constraint: core systems, workflows and controls were not designed for AI participation. As a result, much of the required context and control is recreated for each deployment.

Across the five capabilities, two requirements recur: banking systems that AI can use with the necessary context, and a control plane that establishes which model or agent is acting, what it may do and what must be recorded. Neither is a use case, so neither is captured by the number of live deployments.

The five priorities in the executive summary follow from this diagnosis. Their sequence should be determined by the least mature capability. The more meaningful measure of progress is not only how many use cases are live, but whether each successive deployment becomes easier and less expensive to deliver.

> The next phase of AI in Indian banking will be defined not by how much AI banks deploy, but by how much AI they can responsibly carry.

Survey & Research Methodology

Research objective and design

Banking Frontiers and Zeta jointly developed this study to examine AI adoption across Indian banks and NBFCs, with particular attention to the enterprise capabilities required to progress from experimentation to scaled deployment.

The research captured CXO perspectives on AI adoption, organisational readiness, data and infrastructure, operations, risk and governance, software development, workforce capability, partnerships and business outcomes.

Respondent profile

The study covered 40 CXO respondents from 18 financial institutions, including large and mid-sized private-sector banks, leading public-sector banks and major NBFCs. Respondents comprised CIOs and CTOs, CDOs, COOs and CROs, providing perspectives across technology, data, operations and risk.

Survey and analysis

Separate structured questionnaires were developed for each CXO persona. These combined quantitative questions on adoption, maturity, constraints and priorities with qualitative questions that provided context for the reported scores. Banking Frontiers conducted respondent outreach and data collection through direct engagement with participating CXOs.

Quantitative responses were analysed for patterns across adoption and readiness. Qualitative responses were reviewed for recurring themes and explanations. Only findings considered sufficiently supported and material were included in the report. The survey evidence was interpreted using Banking Frontiers' industry perspective and Zeta's banking technology expertise.

AI tools assisted with aspects of analysis and interpretation. Human oversight was maintained throughout research design, data collection, analysis, validation and report development.

Interpreting the findings

The study provides a directional view of AI adoption and readiness among participating Indian banks and NBFCs. It is not intended to represent the industry statistically.

Respondent bases vary across personas and questions and become smaller when segmented by institution type or functional group. Percentages, averages and comparisons should therefore be interpreted as directional indicators rather than definitive industry-wide conclusions.

Acknowledgements

Banking Frontiers and Zeta thank the participating CIOs, CTOs, CDOs, COOs and CROs for contributing their time and perspectives to this research.

About Us

Zeta is a next-generation banking technology company delivering cloud-native platforms and platform-led solutions across core banking, cards, payments, lending, identity and digital banking.

Zeta has 1700+ employees with over 70% in technology roles across locations in the US, Middle East, and Asia – representing one of the largest and most capable teams ever assembled in banking tech.

Visit us at

www.zeta.tech

Talk to our
experts today.