

Ambition is everywhere. Readiness isn't.

A six-dimension read on where enterprises — and Indian BFSI — actually stand on AI, and why the gap is structure, not ambition.

Edition v1 · Baseline · India & BFSI focus

About this edition. A v1 baseline — Agentora's six-dimension readiness framework applied to the best available public research (McKinsey, Gartner, S&P Global, NASSCOM, RBI; cited in Sources). From v2, figures are recalibrated with anonymized, aggregated data from real Discovery Workshops — only counts and averages aggregate; names and evidence never leave the engagement.

01

Executive summary

Across enterprise India and the world, the gap in AI is no longer *ambition* — it is *structure*. Boards have mandated AI; budgets exist; pilots are everywhere. Adoption is near-universal; readiness is not.

88%

use AI in ≥1 business function

up from 55% in 2023 —
McKinsey¹

>80%

see no material EBIT impact from gen-AI

McKinsey¹

~7%

have scaled gen-AI enterprise-wide

McKinsey¹

12%

say their data is genuinely AI-ready

Precisely / LeBow⁴

1. **Enthusiasm outruns readiness.** Near-universal adoption, almost no value capture. Gartner expected **30% of gen-AI projects abandoned after POC by end-2025** and finds **85% fail on poor data quality**.²
2. **Governance is the universal bottleneck — worst where stakes are highest.** Despite 88% AI usage, comprehensive governance frameworks sit in the low single-to-double digits of firms, and **60% of legal/compliance/audit leaders now name AI their top risk**.³
3. **Data readiness is over-reported.** Only 12% call their data AI-ready; **~80% say AI is held back by data-access challenges**.⁴ Storage is solved; accessibility and lineage are not.
4. **Almost no one measures knowledge maturity.** The ability to capture knowledge in a reusable, traceable form is the least-developed dimension — and the strongest differentiator of the rare "Leading" organization.

The takeaway: most enterprises don't have an AI problem. They have a structure problem wearing an AI costume — and structure is fixable in weeks, not quarters.

The readiness model


Six dimensions, each scored 0–100, rolled into an Overall Readiness Score and a maturity band. They mirror Agentora’s live assessment, so you can move straight from this report to your own score.

#	DIMENSION	THE QUESTION IT ANSWERS
1	Strategy & business clarity	Are AI use cases tied to a measurable business outcome?
2	Data & infrastructure	Is data quality, accessibility, and lineage AI-ready?
3	Technology & architecture	Cloud, LLM adoption, and MLOps/observability maturity?
4	People & process	Ownership, AI literacy, and change control in place?
5	Governance & risk	Explainability, model review/approval, classification, residency?
6	Knowledge maturity	Is organizational knowledge captured, standardized, traceable?

Maturity bands: Foundational (0–34) · Emerging (35–64) · Advancing (65–84) · Leading (85–100).

The overall picture: tall on intent, short on control

The *shape* of the average enterprise profile matters more than any single number.



Bar lengths illustrate the characteristic profile shape; callouts are the sourced public proxies behind each. First-party distributions replace these in v2.

Momentum concentrates in the first dimension and collapses across the last two. S&P Global found the average organization **scrapped 46% of its AI POCs before production**, and MIT estimates **~80% of pilot-to-production work is data, governance, and integration** — not modeling.²

Readiness by dimension

Strategy & business clarity — usually strongest, often falsely so

Most enterprises have an AI *ambition*; fewer have a *ranked* set of use cases. With **80% reporting 50+ gen-AI use cases in the pipeline but only a handful in production**³, intent without prioritization looks like readiness and behaves like drift.

Data & infrastructure — the most over-rated dimension

Only **12% call their data AI-ready**, **51% now cite data governance as a top data-integrity challenge (~89% YoY jump)**, and **~80% say data access constrains their AI**.⁴ Teams discover this only when a use case needs clean, joined, permissioned data — after the pilot is greenlit.

Technology & architecture — capable, shallow on operations

Cloud and LLM access are widespread; production MLOps, evaluation, and observability are rare — reflected in the **~30–46% of projects scrapped between POC and production**.²

People & process — the quiet blocker

AI literacy, clear model ownership, and change control lag the technology; the **skills gap is repeatedly cited as the single biggest barrier**.⁴ Without them, even a good model has no accountable home.

Governance & risk — the universal bottleneck

44% say their governance process is too slow; 56% say intake-to-production takes 6–18 months; many claim a framework, but **fewer than a quarter have fully implemented the controls**.³ See §06.

Knowledge maturity — the invisible dimension

Almost no one measures whether the knowledge an AI system depends on is captured and traceable. Least-developed, clearest marker of a "Leading" organization. See §07.

Industry breakdowns

INDUSTRY	STRENGTH	GAP	POSTURE
Banking (BFSI)	Data infra, risk culture	Explainability & audit at model level	Held back by governance, not capability
Insurance	Actuarial/data depth	Legacy systems, transparency	Emerging; claims & underwriting lead
Healthcare	Motivated use cases	Data governance, safety assurance	Cautious; governance-gated
Manufacturing	Operational data (IoT/OT)	Data accessibility, talent	Emerging; opportunity-rich
Retail / CPG	Customer data, fast iteration	Fragmented data, governance	Advancing on experimentation
Logistics	Optimization use cases	Integration, data quality	Emerging
Government	Mandate, scale	Procurement, compliance, talent	Foundational; compliance-first
IT services	Talent, delivery	Reusable methodology across clients	Advancing; a channel for readiness

The India picture. India's AI market is projected to reach **\$17B by 2027 (25–35% CAGR)**, with BFSI, retail/CPG, healthcare and industrials driving **~60% of the net-new value add**.⁵ Intent is high — **27% of Indian companies already report AI agents in production or at scale** — but production maturity trails intent.⁵

Beachhead — BFSI. Real budgets, dense regulation (RBI, IRDAI, DPDP Act), low tolerance for unexplained decisions. Readiness is gated by governance, not capability: **67% of financial institutions want to explore AI, but only ~21% have deployed it in production**.⁶

The governance gap

In a regulated enterprise, a use case that cannot be explained is not a use case — it is a liability with a demo attached. Governance fails most often on four checks:

1. **Explainability** — can the decision be explained to a regulator, after the fact?
2. **Audit trail** — an unbroken record from input → model → output → the human who signed off?
3. **Accountable human** — a named owner on the record, not "the model decided"?
4. **Classification & residency** — is sensitive data classified, controlled, and resident where it must be?

The regulatory signal is now explicit. On 13 August 2025 the RBI released its **FREE-AI Committee report** — a Framework for Responsible and Ethical Enablement of AI for banks, NBFCs, and payment operators: 7 Sutras, 6 pillars, 26 recommendations, expected to be operationalized via master directions over 12–24 months.⁶ With the DPDP Act's consent and residency duties, the direction is unambiguous — responsible-AI governance is moving from optional to expected.

The organizations treating governance as a **day-one architecture decision** — not a bolt-on — are the ones whose pilots actually reach production.

Knowledge maturity — the differentiator no one measures

AI is only as good as the organizational knowledge it can draw on. Yet "is our knowledge captured, standardized, and traceable?" is the question enterprises are least able to answer — and the one almost no survey even asks.

For AI quality

Ungrounded models hallucinate; grounded ones cite. Retrieval is only as good as the knowledge beneath it — which is why **~80% of pilot-to-production effort is data and knowledge engineering**.²

For compounding returns

Structure knowledge from each initiative and the next one is faster. Skip it and you restart at zero every time — the trap that makes traditional consulting so expensive.

The rare "Leading" organizations treat knowledge as an asset to be structured, not a byproduct to be filed. It is the strongest single predictor of durable readiness — and the emptiest column on most scorecards.

Four readiness archetypes

The Enthusiast Without Foundations

High strategy, low data/governance. Lots of pilots, little production.

Needs: structure & sequencing

The Compliance-First Cautious

Strong governance instinct, slow to start. Waiting for a certainty that never comes.

Needs: a bounded, low-risk first win

The Capable but Fragmented

Good tech and data in silos, no shared method. Wheels reinvented per team.

Needs: reusable methodology + knowledge layer

The Quiet Leader

Balanced across all six, including knowledge. Rare. Ships and governs.

Needs: keep compounding, don't slow down

What the leaders do differently

1. **Score use cases before funding them** — $\text{impact} \times \text{feasibility} \div \text{governance risk}$, ranked, on the record.
2. **Put governance on day one** — explainability and audit trails are architecture, not paperwork.
3. **Name an accountable human for every model** — owner and steward, before deployment.
4. **Capture knowledge as they go** — each initiative structured so the next is faster.
5. **Choose architecture on merit, not vendor pull** — recommendations follow the problem, not a referral fee.

None are technology moves. All are structural — which is why they are learnable and fast. The 11% of agentic pilots that *do* reach production reportedly return **171% ROI**.²

Close your gap

1. **Measure honestly** — free AI Readiness Assessment, six-dimension score → agentora.in/assessment
2. **Score your backlog** — the $(\text{impact} \times \text{feasibility}) \div \text{governance-risk}$ template → agentora.in/resources
3. **Fix the weakest axis first** — usually governance or knowledge, not the model.
4. **Prove it on one bounded engagement** — a fixed-price Discovery Workshop turns the score into a roadmap in days → agentora.in

Sources

1. McKinsey & Company — *The State of AI 2025* (88% adoption; >80% no material EBIT impact; ~7% enterprise-wide scaling).
2. Gartner (30% of gen-AI projects abandoned after POC by end-2025; 85% fail on data quality; 89% of agentic pilots don't reach production; 11% return 171% ROI), S&P Global Market Intelligence 2025 (46% of POCs scrapped), MIT (~80% of pilot-to-production work is data/governance/integration).
3. ModelOp *2025 AI Governance Benchmark* & Thomson Reuters Institute (minority of firms with comprehensive governance; 60% of legal/compliance leaders cite AI as top risk; 44% governance too slow; 56% 6–18-month intake-to-production; <25% controls implemented).
4. Precisely / Drexel LeBow (12% AI-ready data; 51% cite data governance, ~89% YoY rise) & Cloudera 2026 (~80% held back by data access).
5. NASSCOM (AI Adoption Index; Enterprise AI-agents 2025) & NASSCOM–BCG (India AI market ~\$17B by 2027; BFSI/retail/healthcare/industrials ≈60% of net-new value add).
6. Reserve Bank of India — *FREE-AI Committee Report* (13 Aug 2025; 7 Sutras, 6 pillars, 26 recommendations) and Indian regulated-entity adoption (67% exploring; ~21% in production).

Figures are third-party market indicators as of mid-2026; verify the latest source text before external publication. © 2026 Agentora Technologies Pvt. Ltd. · This v1 baseline becomes a first-party benchmark from v2, calibrated on anonymized engagement data.