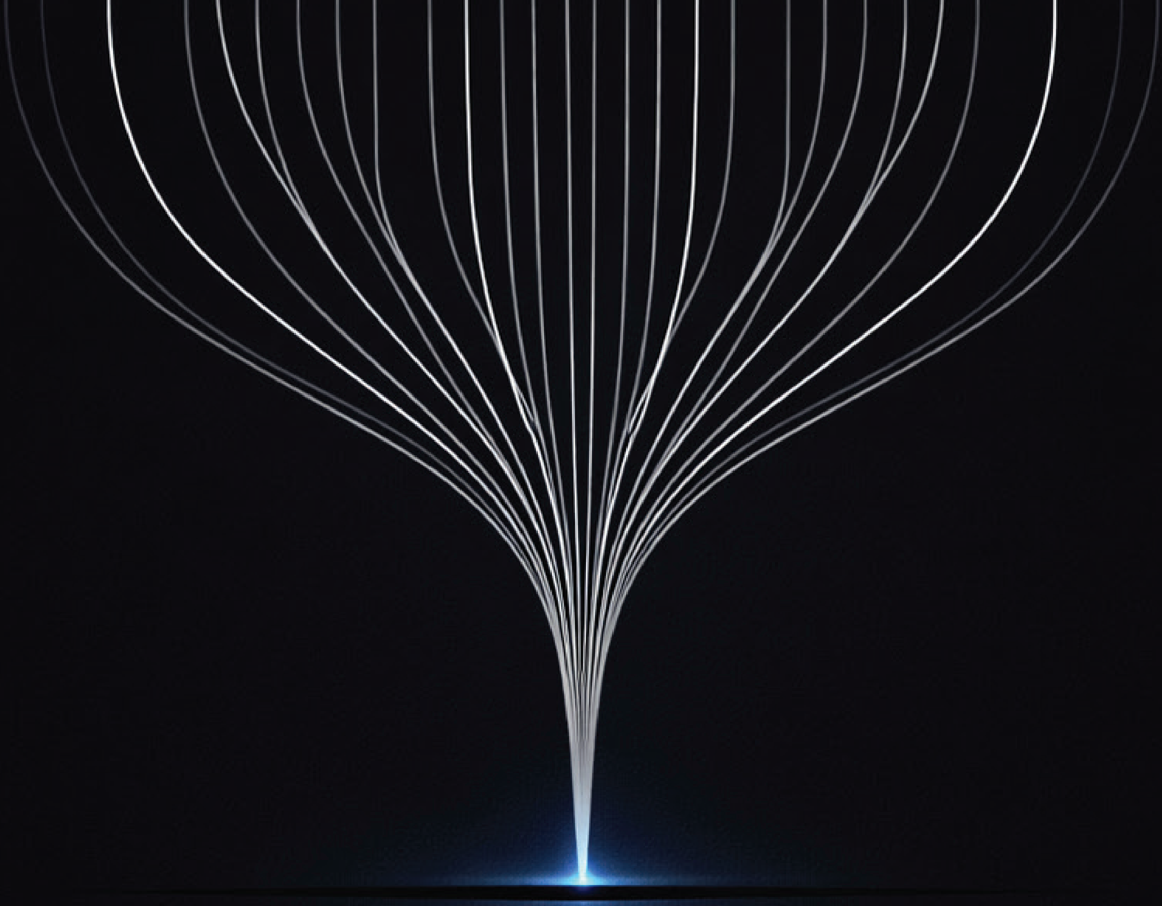

THE DECISION BOTTLENECK NO ONE TALKS ABOUT

GLIB

A GLIB Research Report on AI in BFSI

Indian BFSI leaders no longer debate whether to use AI

That decision already sits behind them.

The real question now is how to scale AI from isolated tasks to governed, autonomous decisions without losing control.

Bridging the gap starts with asking the right questions. At GLIB, we've observed this firsth and across 50+ BFSI implementation from lending automation at a top African bank to document orchestration for India's largest life insurer where siloed AI tools throttled holistic decisions until agentic coordination unlocked scale. This report unpacks the five questions that determine success or failure in agentic AI adoption.

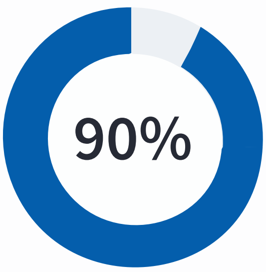

Indian BFSI firms have AI in their strategy.

Strategic intent for AI is nearly universal across the sector.

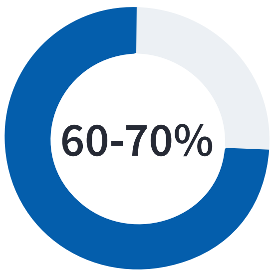

are using AI in core functions.

AI has moved beyond planning and into active operational use.

1. AUTHORITY & ACCOUNTABILITY:
If an AI agent makes a lending decision, who is liable if it fails? How do we maintain regulatory accountability for autonomous decisions?

2. DATA GOVERNANCE & QUALITY:
How do we ensure 99% accuracy in structured data extraction before we feed agents unreliable inputs?

3. CHANGE MANAGEMENT & RESISTANCE:
How do we retrain 40-50% of our workforce when 50% of roles get reshaped? What's the timeline?

4. CAPITAL ALLOCATION & ROI:
What's the total cost of orchestration (infra, talent, governance) vs. gains from autonomous decisions? What's the payback period?

5. COMPETITIVE SPEED vs. RISK:
If we move more slowly on agentic AI for safety, will competitors with aggressive approaches capture market share?

We have celebrated the adoption of AI tools, but have overlooked the integration of AI-driven decisions. This creates a patchwork of isolated efficiencies rather than a cohesive, intelligent system. The gap does not lie in technology. It sits in decisions.

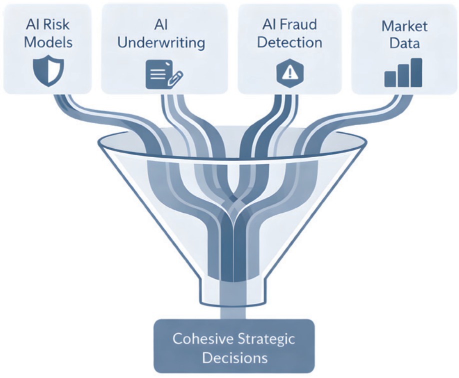

The potential of multiple AI tools is throttled by a lack of coordination across critical decision flows. Insights live in silos, but financial institutions' ability to make fast, holistic decisions remains constrained.

Fragmented decision-making, even when powered by advanced AI, leads to:
- Slower market response
- Sub-optimal capital allocation
- Missed synergies
- Escalating complexity

This report shows how GLIB's 4-layer agentic framework answers each question through proven deployments at Absa Bank, SBI Life, and global ratings agencies.

The Gap Isn't Technology. It's Orchestration.

The barrier to scale isn't the quality of the models or the power of the platforms. It is the absence of a framework to manage and govern autonomous systems—the lack of orchestration.

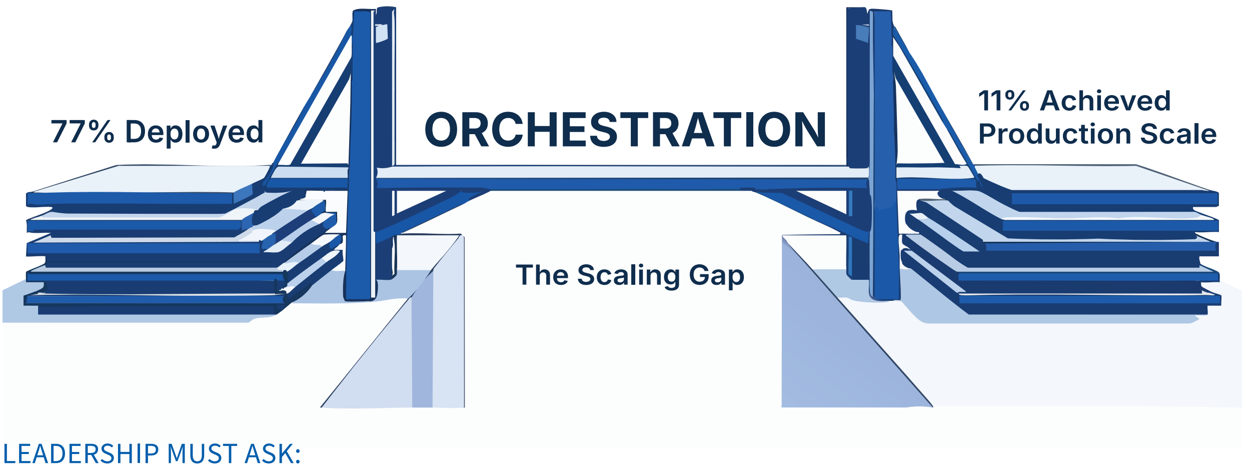

LEADERSHIP MUST ASK:
1. Where does AI influence decision authority vs. task automation?
2. What's the cost of fragmented decision flows?
3. Who or what decides, and how consistently?
4. How do we benchmark against market leaders to uncover missing efficiencies in agentic orchestration?

The Hidden Cost of Fragmented Intelligence

AI adoption in **BFSI** is accelerating: 77% of global banks deployed AI by 2025, and 74% of Indian firms piloted GenAI, but true system-level integration remains rare, with only 11% reaching production scale across departments.

Credit teams leverage AI for eligibility scoring, risk for anomaly detection, yet without orchestration, outputs clash: a "low-risk" profile from credit might ignore fund trail signals flagged by risk. This disconnect forces manual reconciliation, where executives rebuild context from disparate dashboards, delaying decisions by days and capping throughput at pre-AI levels.

When intelligence doesn't connect, organizations can't build decision frameworks. Risk teams see one version of reality while credit teams see another.

Decisions slow down where coordination matters most. When manual escalation becomes a safety net, senior teams intervene much later.

When fragmented intelligence forces humans to act as integrators, it keeps institutions safe, but it also caps scale.

The Production Paradox

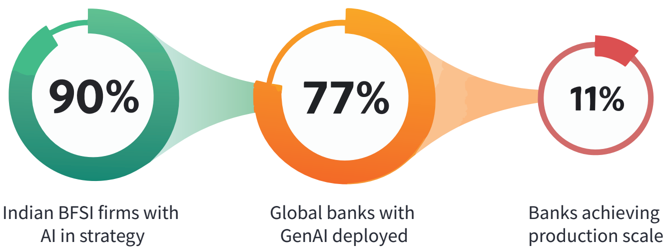

77% of global banks have deployed GenAI, yet only 11% have achieved production scale.

Agentic AI Maturity Across BFSI Decision Domains

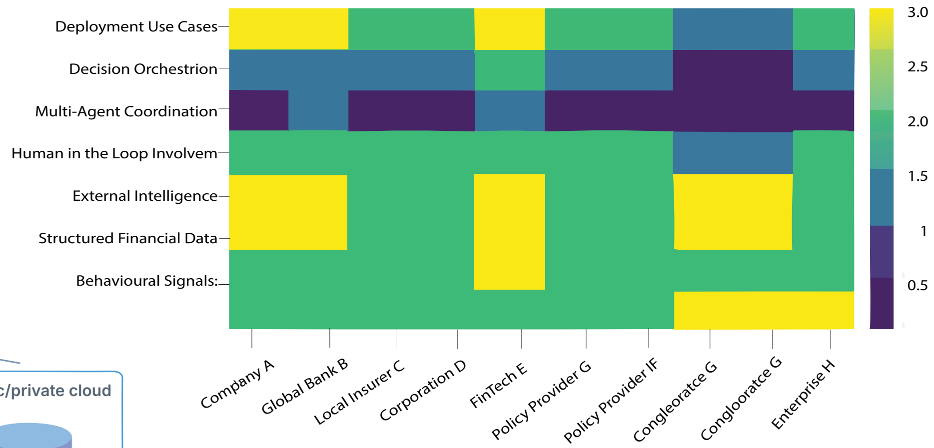

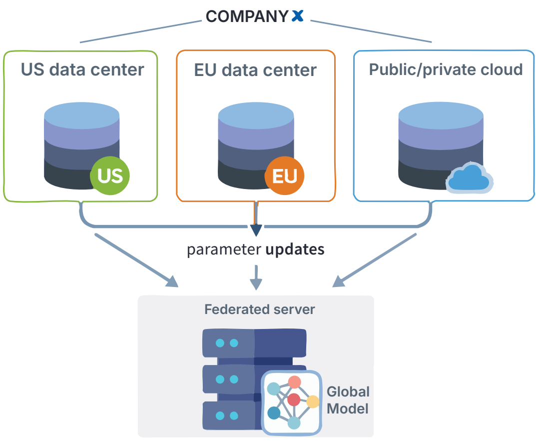

AI initiatives must create collective intelligence that compounds as organizations scale.

Question to ask: Where does intelligence stop flowing across teams in your decision chain?

Documents Are Not Paperwork.

When documents are treated as paperwork, their signals get flattened. Credit teams focus on eligibility. Risk teams look for anomalies. Compliance teams check sufficiency. Each function touches the same documents, but no single system carries their combined meaning forward.

How Documents Lose Meaning Across Teams

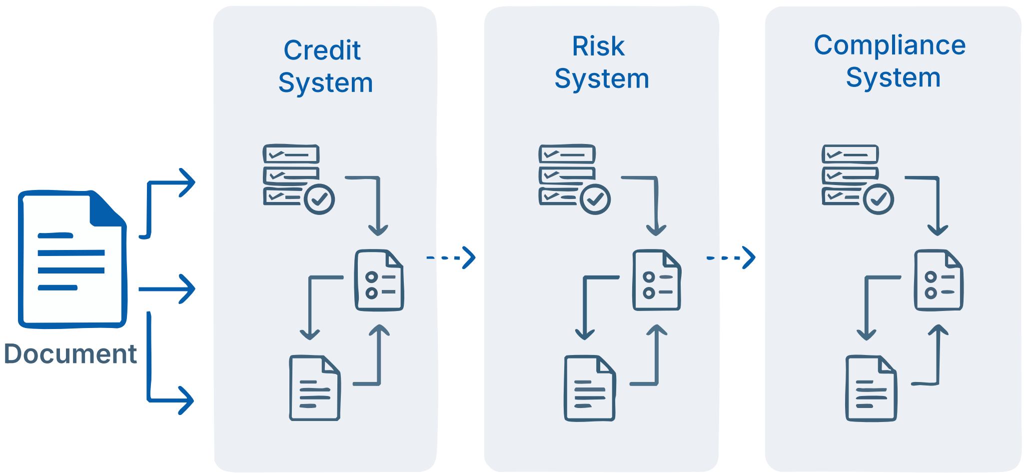

Parallel interpretation without shared memory

RBI notes the rapid adoption of AI-related systems in bank reporting, signalling increasing reliance on intelligent data processing.

Documents are behavior carriers. They reflect liabilities, intent, and risk, not just content to automate.

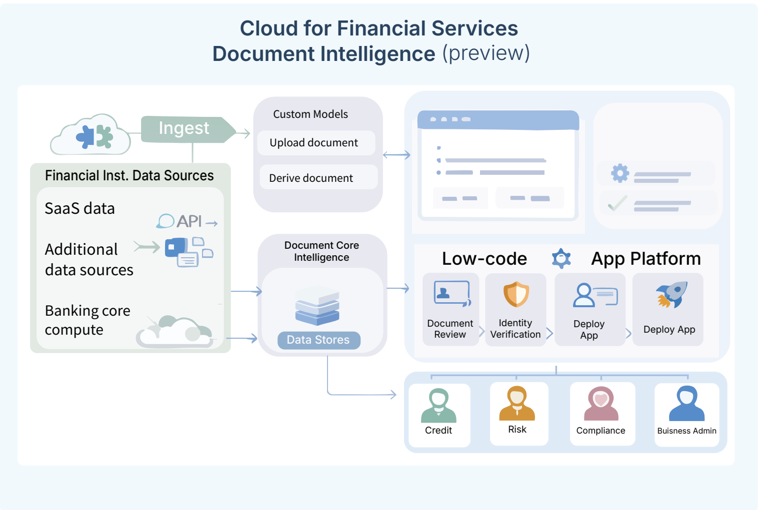

Speed improves when documents move faster. Decisions improve only when the document's meaning carries forward.

Why Automation Plateaued in **BFSI**

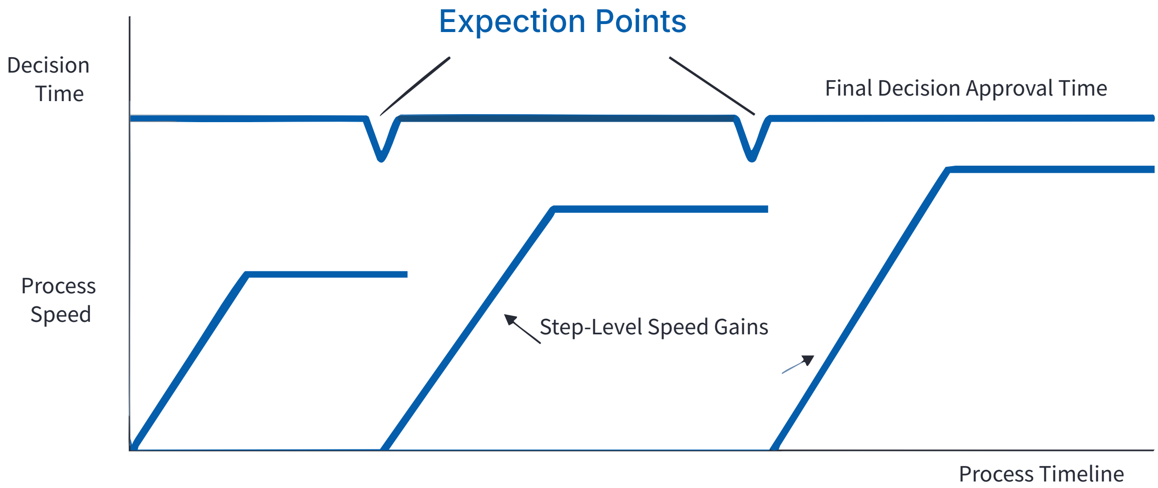

By **2025**, **77%** of global banks* will have deployed AI in some form
Despite **74%** of Indian BFSI initiating GenAI PoCs, only **11%** have achieved production scale.* This reveals an orchestration gap between capability and deployment.

PoCs succeed in silos (**e.g.**, **chatbots, fraud pilots**) but falter at scale due to regulatory hurdles (RBI data localization/PII redaction), legacy integration barriers, data quality issues, and skyrocketing infra costs—leaving **89%** stuck without cross-department coordination. Agentic systems bridge this by governing multi-tool workflows under supervision.

The next constraint is not the speed of individual tasks. It's the coordination of judgment across functions. In lending, credit AI greenlights eligibility while risk flags fund trail anomalies; without orchestration, compliance vetoes late due to **PII silos**.

Legacy systems exacerbate this:
disparate outputs force humans to aggregate insights manually, eroding context and inviting bias.

Solving this requires systems that can reason, remember, and escalate deliberately, not just execute.

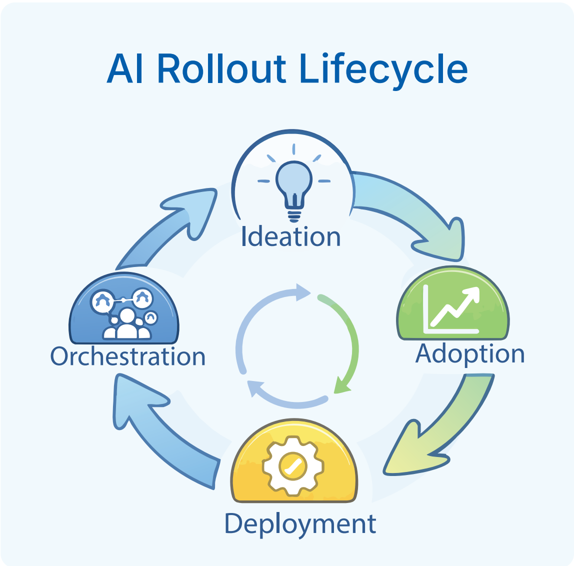

The AI Evolution:
From Text Prediction to Autonomous Judgment

AI has layered rapidly.

1. Large Language Models (LLMs) power prediction, turning vast data into coherent text/code.
2. Generative AI (GenAI) builds creative outputs like reports/summaries.
3. AI Agents add execution via tools (**e.g., API calls**).
4. Agentic AI orchestrates it all: setting goals, planning steps, adapting in loops, governing sub-agents for complex workflows.

service costs, underwriting copilots cutting loan TAT via alt-data scoring, and ops automation reducing manual processing to 1/10th costs (**e.g., CRM/loan origination**). NBFCs lead in BI insights; banks scale cybersecurity/customer care. Yet gains cap without governance.

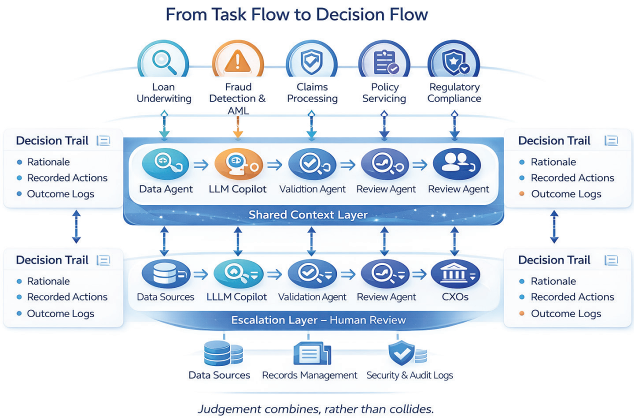

Layer	Core Capability	BFSI Example	Key Limit
LLM	Text prediction	Chat responses	No actions/memory
GenAI	Content generation	Risk summaries	Reactive, no orchestration
Agents	Tool execution	Data fetch	Fixed tasks
Agentic	Goal-base autonomy	Cross-team decisions	None- scales judgment

Generative AI Powerful, But Not a Decision System. GenAI could boost Indian banking operational efficiency by up to ~**46%**.*
How? Through AI chatbots slashing customer

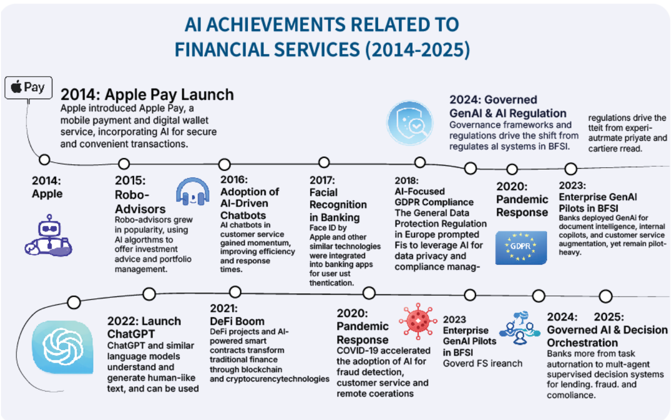

Used well, it removes friction from many knowledge-heavy tasks. Generative systems produce outputs, they don't inherently govern cross-functional risk and accountability. Authority flows are still human-centred. Generative AI excels at producing content. Decision systems require coordination, memory, and supervision across time.

Without these elements, intelligence remains reactive. It informs judgment but does not structure it. Understanding this boundary is essential before attempting to scale AI deeper into **BFSI workflows**.

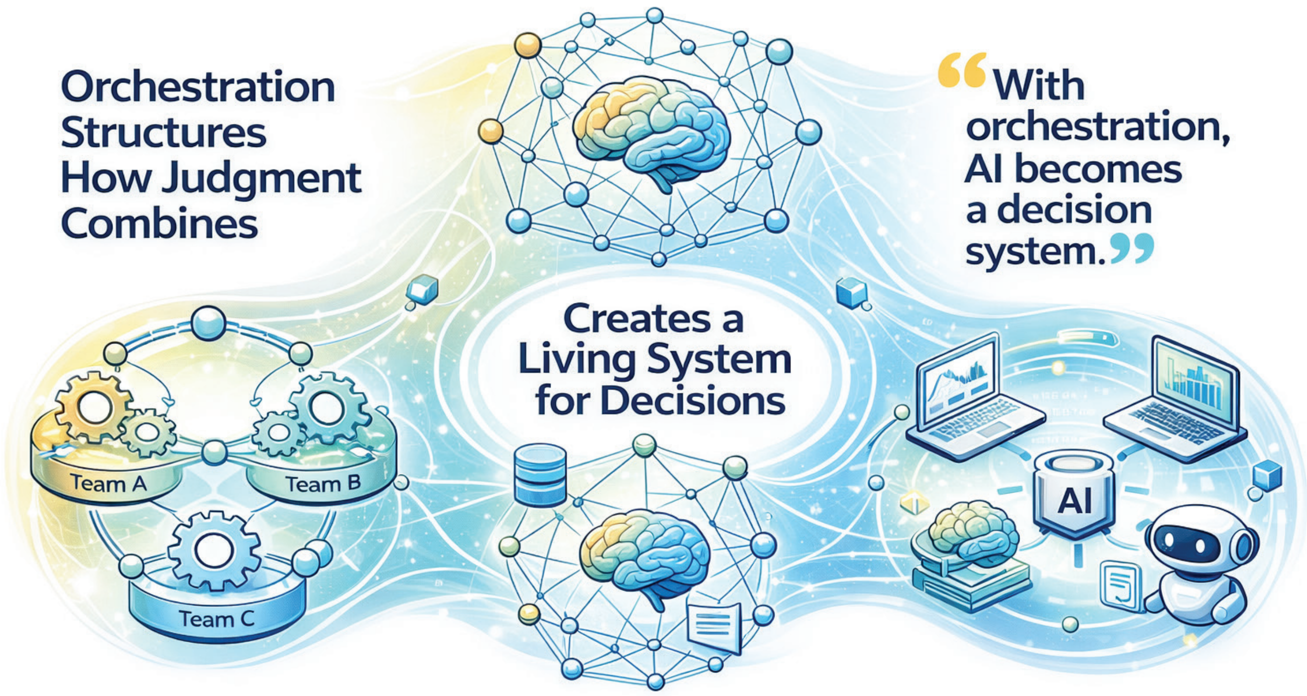

The Missing Layer: Orchestration of Decisions

The global AI in banking market is projected to grow to ~USD 379 billion by 2034 (30.6% CAGR).*

With AI implementation, breakdowns occur between teams, not during individual tasks like data entry. Without a connecting layer, each team’s judgment and context is lost at handoffs.

Orchestration structures how judgment combines. It defines when a decision proceeds, pauses, or escalates to the right person. It preserves context across steps, records rationale, and enables proactive supervision.

AI orchestration transforms a collection of isolated tools into an intelligent, connected system.

Enter GLIB framework...
GLIB's 4-Layer Agentic Framework: The Missing Orchestration Layer in Action

Enter GLIB framework...

GLIB's **4-Layer** Agentic Framework:The Missing Orchestration Layer in Action

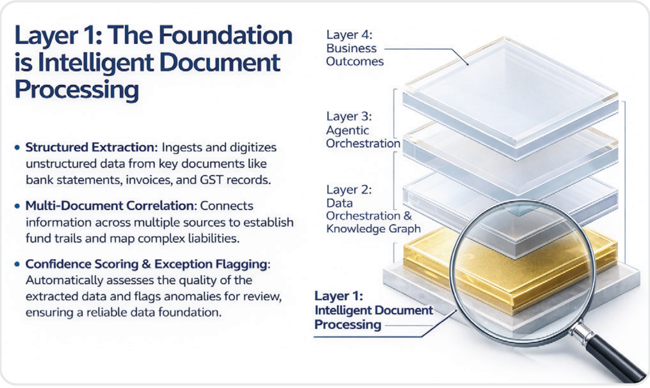

Layer 1 (IDP Foundation) delivers structured truth from unstructured docs with 99% accurate extraction of bank statements, invoices, financials .

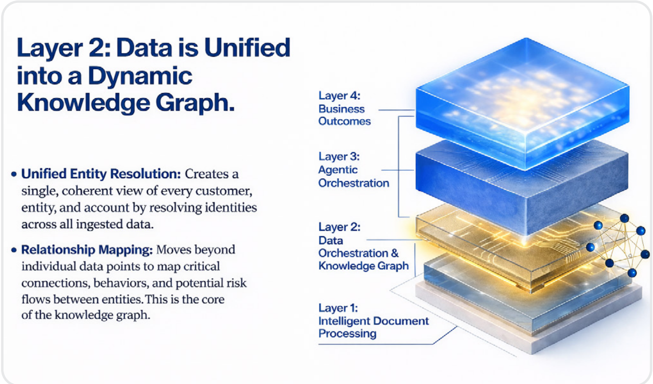

Layer 2 (Knowledge Graph) unifies insights into dynamic graphs so that fund trails surface behavior patterns.

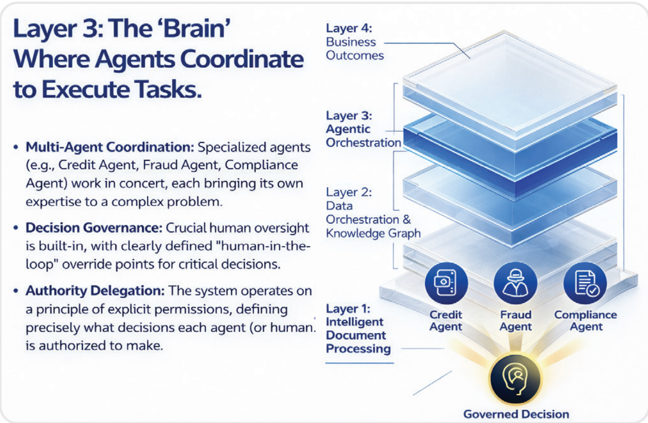

Layer 3 (Agentic Brain) coordinates specialized agents for judgment: reasoning, memory, deliberate escalation with human-in-loop governance.

Layer 4 (Business Outcomes) scales to ROI.

Glib transforms silos into governed, compounding intelligence, the orchestration that BFSI scale demands.

What Agentic AI Really Means in BFSI

AI agents accounted for **17% of total AI value in 2025** across industries.*

BFSI agentic AI systems are not designed to replace human judgment. They delegate judgment safely, within defined boundaries. It behaves autonomously as a specialized decision participant.

Structure without delegation limits scale. Agentic AI sits between these extremes. It coordinates how decisions form, evolve, and conclude. It remembers prior outcomes. It tracks unresolved risk. It escalates deliberately, not reactively.

Most importantly, agentic systems operate under supervision. Human authority remains intact. What changes is the burden. Humans no longer reconstruct context from scratch. They intervene where judgment truly matters.
This is why agentic AI represents a different category from task automation or generative assistance. It does not focus on producing outputs. It focuses on managing decisions as living processes.

For BFSI leaders, the question is not whether to allow autonomy.
It is whether decision systems should remain informal or become explicit and governed. Agentic AI makes that governance possible at scale.

What it Is vs. What it Is Not

Agentic AI Is	It Is Not
Role-defined agents	Unsupervised autonomy
Bounded authority	Black-box decisions
Explicit escalation	Model-driven improvisation

Delegation with Supervision

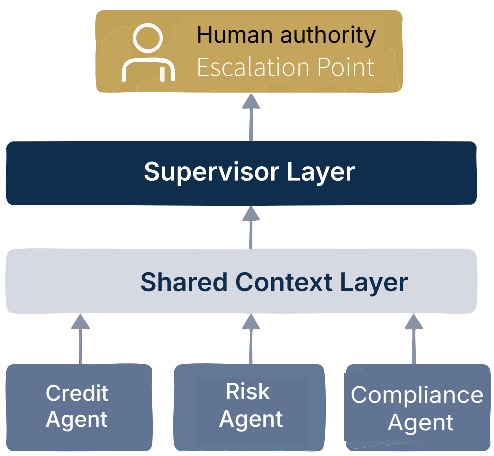

From Linear Workflows to Living Decision Systems

AI could reshape 35–50% of Indian banking roles by 2025.* Most **BFSI** workflows still assume linearity. Information arrives. Checks run. Decisions conclude.

This model worked when volumes were predictable, and exceptions were rare. It struggles now. Today's decisions unfold under constant change. Customer behaviour shifts. Risk signals evolve mid-process. Regulatory interpretation tightens without notice. A workflow designed as a straight line becomes brittle when reality becomes different.

Living decision systems respond differently. They observe continuously. They retain memory across steps. They adapt when conditions change, without tearing down the entire flow. Decisions evolve as signals evolve.
This is the shift that agentic systems enable.

Rather than pushing decisions through predefined paths, organisations allow decisions to **remain open** until confidence stabilises. Agents monitor, re-evaluate, and escalate when thresholds change. Humans intervene with full context.

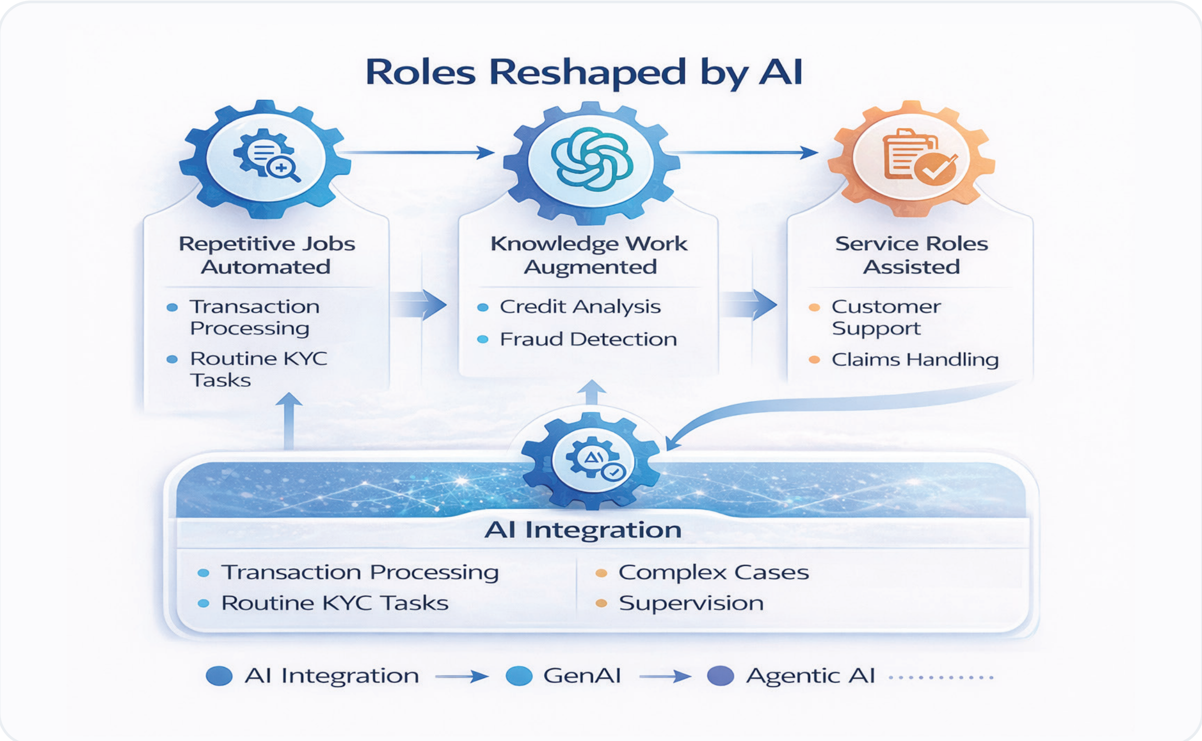

Why Skipping the IDP Layer Breaks

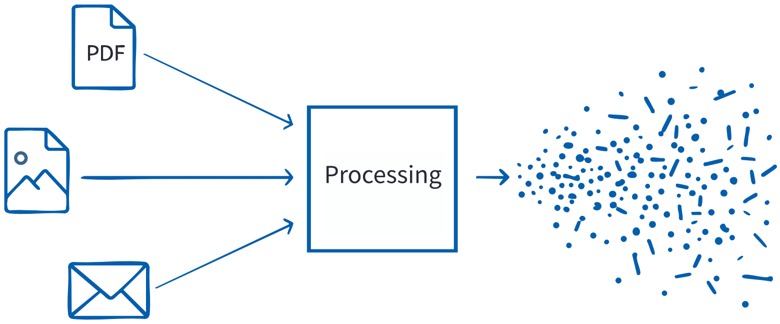

A Teamless Digital survey shows the BFSI industry leads with a 68% AI adoption rate in 2024. Moreover, the **BFSI** industry holds a 30% AI market share.*

AI intelligence systems rest on a simple assumption: The information it receives is reliable. But these critical signals arrive in inconsistent formats like bank statements, invoices, and compliance records.

With partial extractions and lost context, this fragmented information
loses its original meaning.

The focus must shift from simple extraction to preserving intent and context. Domain-aware logic is necessary to interpret financial and regulatory meaning, validate consistency against known rules, and standardize outputs into structured truth before decisions are formed.

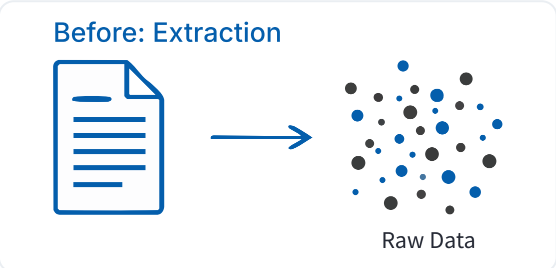

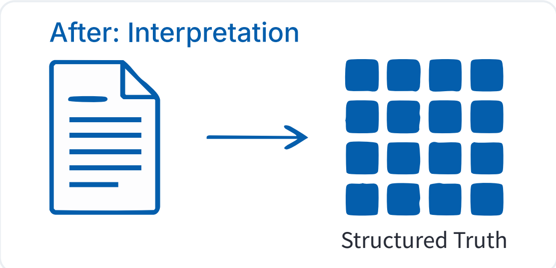

Intelligent document processing establishes a stable reference point for reasoning

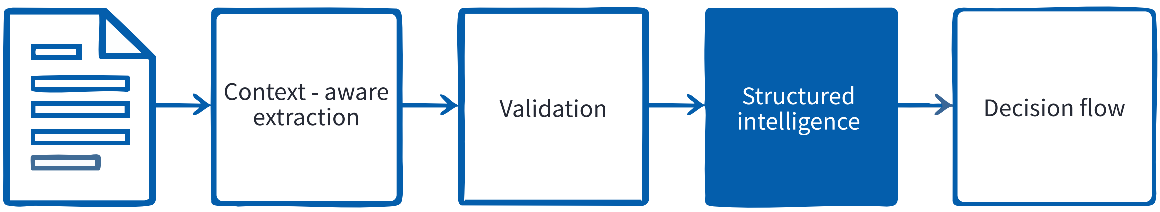

Intelligent Document Processing as the Foundation

Why Decision Systems Depend on Structured Truth?

Indian AI in the **BFSI** market is expected to grow from USD 830 million (2024) to USD 8. 090 billion (2033) (28.8% CAGR).*
Scale without structure amplifies noise. Structured extraction is the necessary baseline for agentic orchestration.

Intelligent Document Processing addresses this problem at its root.

IDP preserves intent and context as documents move through workflows. It applies domain-aware logic to interpret financial meaning, validate consistency, and standardise outputs before decisions form.

With structured, validated document intelligence in place, decision systems stabilize. Context carries forward, and as a result, judgment compounds.
In agentic environments, IDP is the ground truth layer.

Creating a Flow from Documents to Structured Truth

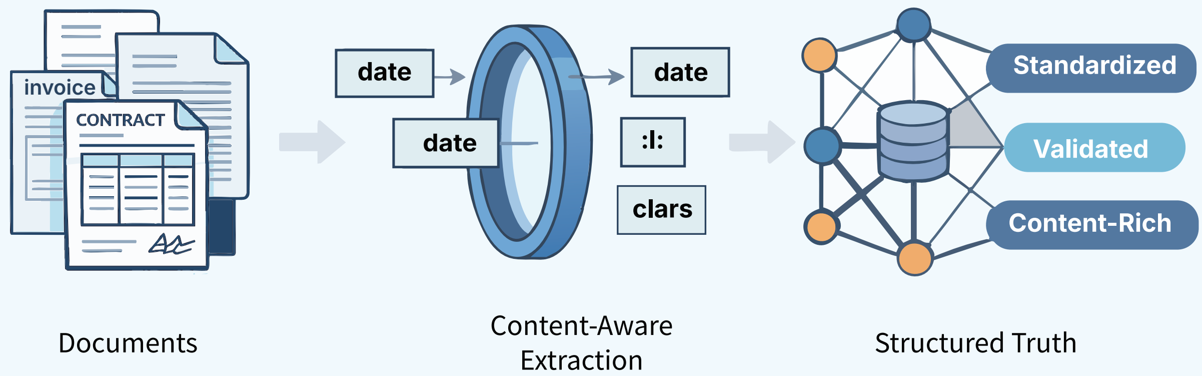

Re-thinking Bank Statement Analysis

A study led by RBI officials* found that the BFSI sector has several use cases for AI, including fraud detection, customer segmentation, and chat automation.
The momentum is led by private banks, and public banks are not far behind.

Most of these AI-powered **BFSI** decisions start with bank statements. Most analysis treats these as static summaries, but don't answer the most pressing question: How does the money actually move?

Behavior lives between transactions, but manual review struggles to surface these patterns at scale. As volume rises, this gap widens, and analysts spend more time per case or escalate uncertain profiles. These slow down decisions and introduce inconsistency.

Systems must look at bank statements through a different lens. Instead of a one-time review, systems must retain memory across periods and accounts.

When bank statements are treated as behavior signals, early warnings surface sooner, and human judgment then focuses on where interpretation truly matters. This is where document intelligence begins to translate into decision intelligence.

Re-thinking Bank Statement Analysis:

From Static Balances to Dynamic Behaviour

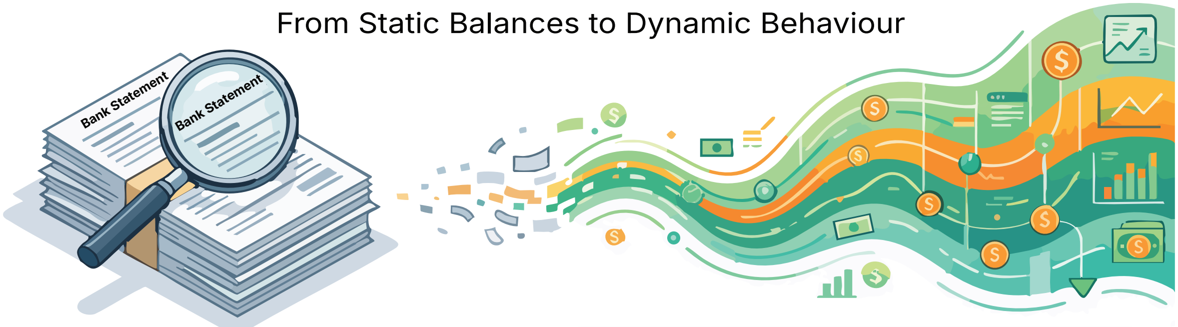

The Old Way: Static Analysis

- **Focuses on Static Snapshots**
  Answers a limited question: "Is there enough money?" by checking balances and averages.
- **Misses Deeper Risk Signals**
  Fails to see emerging pattems, leading to slow, inconsistent, and subjective decisions.
- **Relies on Manual Review**
  Analysts spend more time on cases or escalate uncertain profiles, creating bottlenecks.

The New Way: Behavioural Analysis

- **Analyzes Dynamic Money Flows**
  Answers the critical question: "How does money actully move?" using transaction patterns.
- **Reveals Stress Before Balances Change**
  Financial stress appears in transaction behaviour long before it impacts totals.
- **Improves Decision Confidence**
  Provides eariler visibility into risk, allowing humans to focus where it matters most

Financial Statements as Decision Graphs

Banking AI market growth validates deeper digital transformation.

Financial statements sit at the centre of most high-value decisions in **BFSI**. They are trusted, familiar, and heavily relied upon. Income statements get reviewed for profitability. Balance sheets get checked for leverage. Cash-flow statements get scanned for liquidity. Each view is valid. None is complete on its own.

Financial reality operates as a system. A change in working capital reshapes cash flow. A shift in revenue quality alters balance-sheet strength. An accounting note can materially change how numbers should be interpreted.

When statements are analysed in isolation, these relationships weaken or disappear.

Viewing financial statements as **decision graphs** changes the approach. Instead of linear checks, systems recognise dependencies. They track how one movement affects another. They preserve links between accounts, notes, and disclosures.

When this structure exists, interpretation becomes more consistent. Risk assessment stabilises. Exceptions narrow. Human judgment moves from reconstruction to evaluation.

THE FRAGMENTED VIEW

Income Statement

Balance Sheet

Cash-Flow Statement

Accuracy without coherence

THE GRAPH VIEW

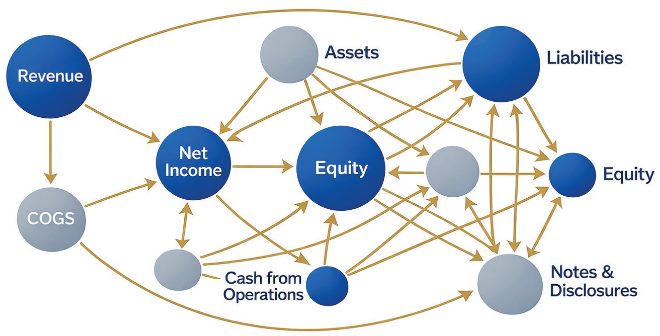

Relationships preserved. Interpretation stabilized

Invoices, GST,and Hidden Liability Risk

Surveys conducted by **RBI** in 2023 and 2024 point out that more than 75%* of banks have deployed AI-powered chatbots for customer service. So, AI is already a part of the decision-making process and powers the credit lifecycle.

Bank financial statement processing is often automated as it's integral to the customer lifecycle. But invoices and tax documents rarely receive the same attention.

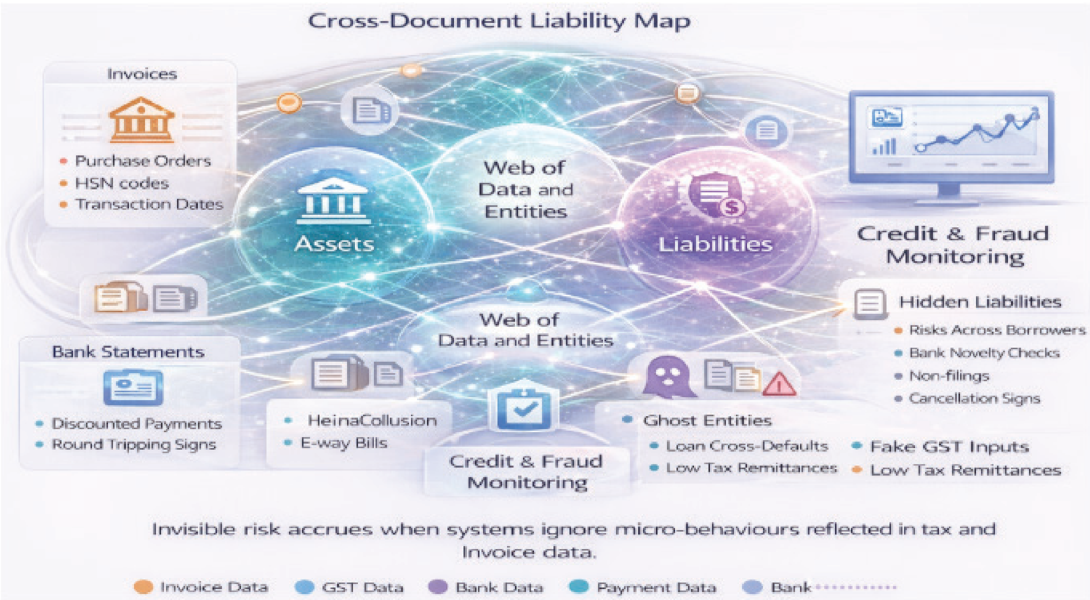

Invoices indicate liabilities in formation. GST filings signal compliance posture in motion. Together, they share exposure long before it appears on balance sheets or risk reports. Ignoring these results in accumulating hidden risks.

Manual reviews struggle to connect invoice reversals, mismatched tax credits, and vendor concentration patterns. This builds exposure slowly without triggering escalation.

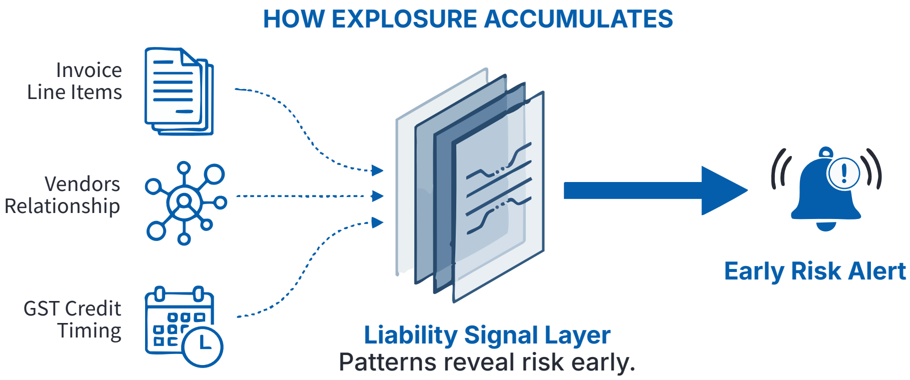

Human-in-the-Loop
Is a Governance Feature

As **AI systems** mature, a familiar question resurfaces.

How much human involvement is enough?

In regulated environments, human oversight is a design requirement. In **BFSI**, accountability sharpens when automation improves.

- Regulators expect clarity on who approved a decision, why it was approved, and what context informed it.
- Customers expect fairness and consistency.

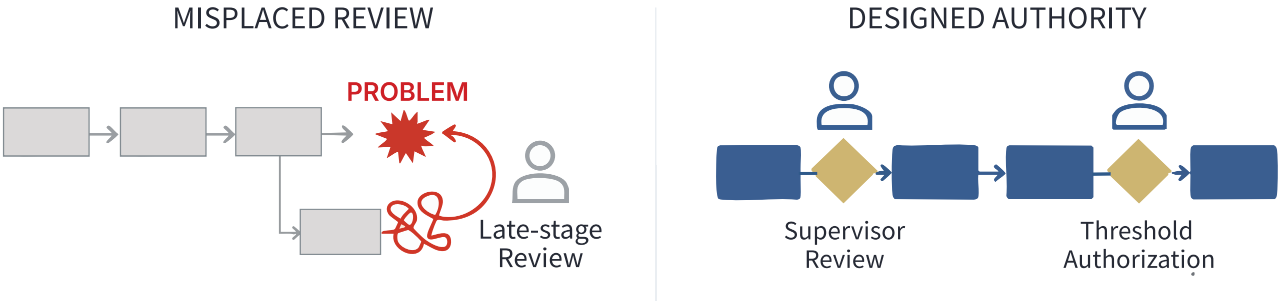

Oversight arrives too late.

Authority placed where it matters

Human-in-the-loop systems meet these expectations when designed deliberately. They intervene at defined control points, authorize thresholds, resolve ambiguity, and supervise patterns instead of transactions.

With agentic AI systems, the role of humans shifts from repetitive validation to accountable judgment.

The design architecture starts by treating documents as control surfaces. IDP establishes structured truth from these inputs. This enables systemic insights like decision graphs and liability intelligence. The entire system operates under a framework of designed authority, ensuring decisions are intelligent and defensible.

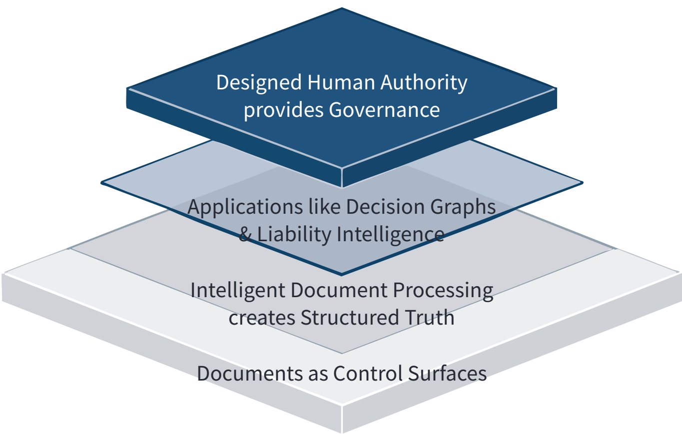

From Validation to Confidence

Validation sits at the heart of every **BFSI workflow**. Data gets checked. Documents get matched. Rules get applied. Yet validation, on its own, rarely creates confidence.

Aspect	Data Validation	Data Verification	Data Validity
Definition	Ensures data meets rules and constraints.	Ensures data handling processes are accurate.	Assesses whether data is accurate and relevant.
Focus	Correctness and completeness.	Accuracy of processing and transformations.	Relevance and representativeness.
When It Happens	At entry or during preparation.	During or after processing.	Continuously throughout the lifecycle.
Examples	Checking required fields, formats, and ranges.	Ensuring aggregation are accurate.	Assessing sampling methods, checking for bias.

Confidence emerges when signals agree, context persists, and rationale remains visible as decisions evolve. Moving from validation to confidence requires a shift in design.

Instead of isolated checks, systems must correlate outcomes. Instead of one-off approvals, reasoning must be preserved across teams. AI agentic systems support real-time explainability by preserving intent and context. When validation moves beyond retrospective audits, confidence builds, and escalation becomes selective.

Trusted AI agents help decisions move with less friction.

AI Auditing Framework

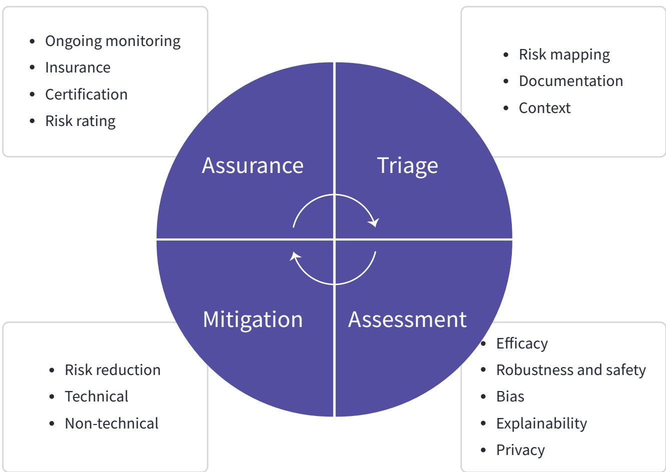

Fund Trails
Reveal Behaviour

AI-based fraud systems could save global banks **£9.6 billion annually by 2026.***

But, more financial controls to check fraud, examine transactions in isolation.
Accounts get checked, parties get verified, and dates get matched separately. However, risk rarely behaves that way. Fraud, stress, and misuse emerge across sequences of movements.

In most cases, singular transactions don't
look suspicious on their own. The underlying fraud pattern only becomes visible when flows are connected.

Traditional monitoring struggles here. Rule-based systems flag known thresholds. Manual review inspects individual entities when relational behavior is what's needed.

Fund trail analysis reframes the problem. It explores how money behaves instead of just verifying whether a transaction is valid.

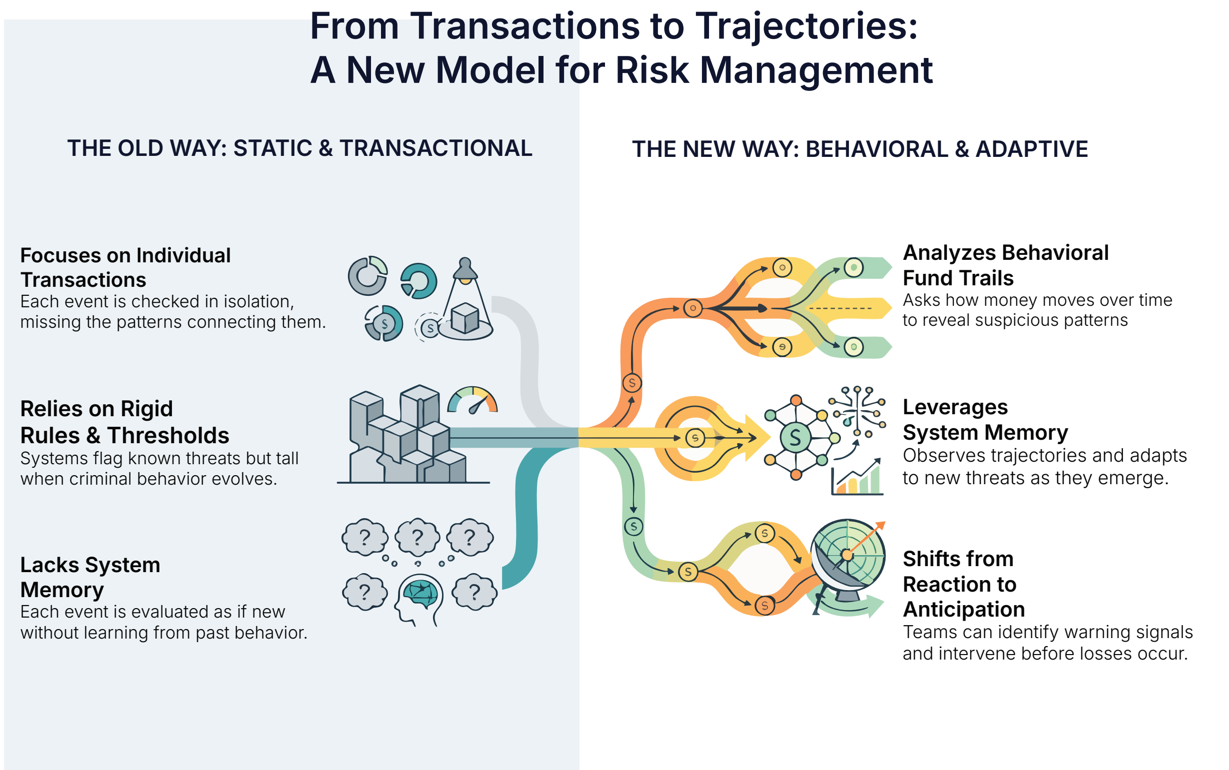

When systems map fund trails, they surface behaviour that static checks overlook. Risk teams can target escalation. This makes human judgment surface earlier with a clearer context. Behavioral fund intelligence shifts risk management from reaction to anticipation.

It reduces dependence on post-event investigation. This is where agentic AI monitoring begins to show its advantage.

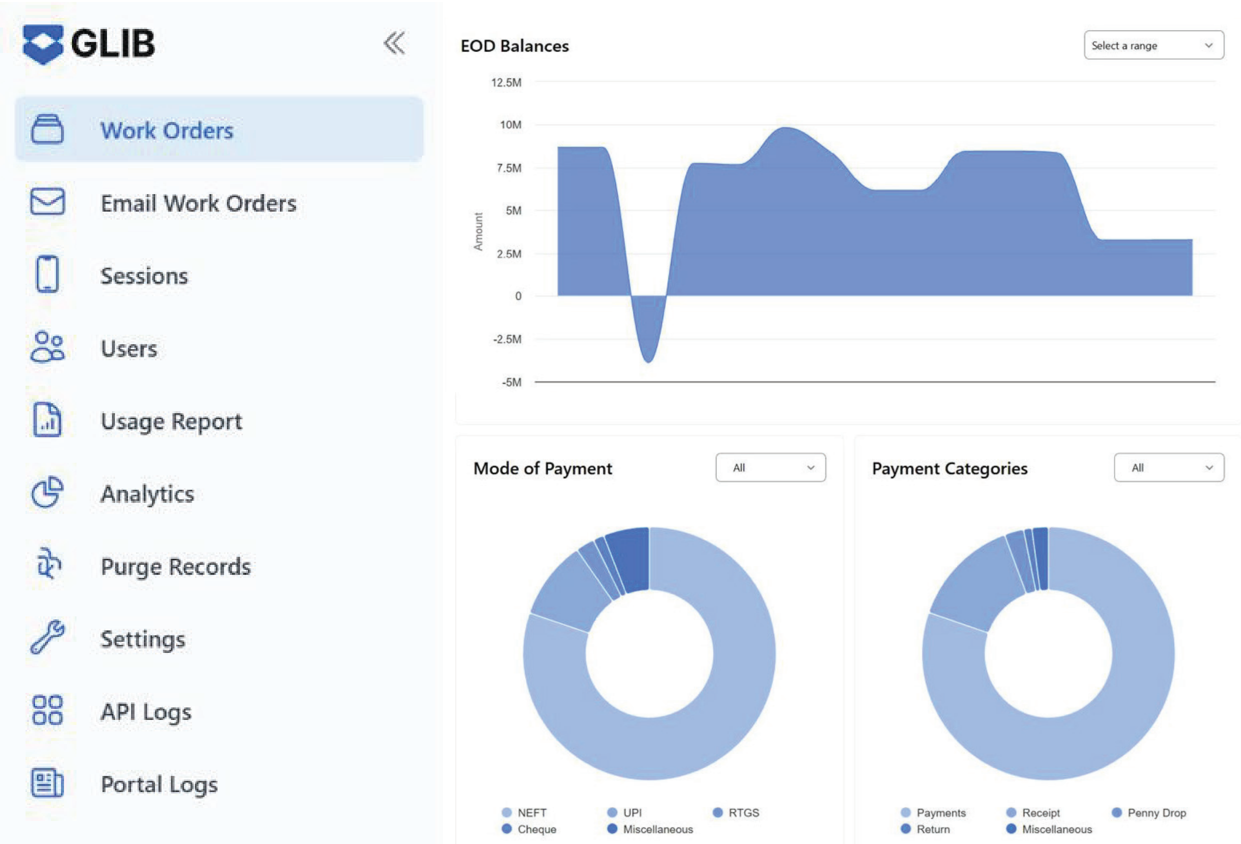

Fraud Needs Memory, Not Just Rules

Fraud evolves faster than rules. Most fraud controls rely on predefined thresholds and known patterns. They work well until behaviour changes. Once it does, systems either over-flag legitimate activity or miss emerging risk entirely. Teams respond by adding more rules. Complexity grows. Signal quality drops.

Static systems evaluate each event as if it were new. They do not remember prior behaviour in context. They do not learn how patterns shift over time. They react repeatedly to symptoms rather than adapting to causes.

Human investigators compensate for this gap. They recall past cases. They recognise subtle deviations. They connect events across accounts and periods. This keeps organisations safe, but it does not scale.

Fraud detection improves when systems gain memory. Agentic AI observes trajectories instead of moments. It enables comparison across time, entities, and behaviours. It distinguishes one-off anomalies from sustained manipulation. It reduces false positives without lowering sensitivity.

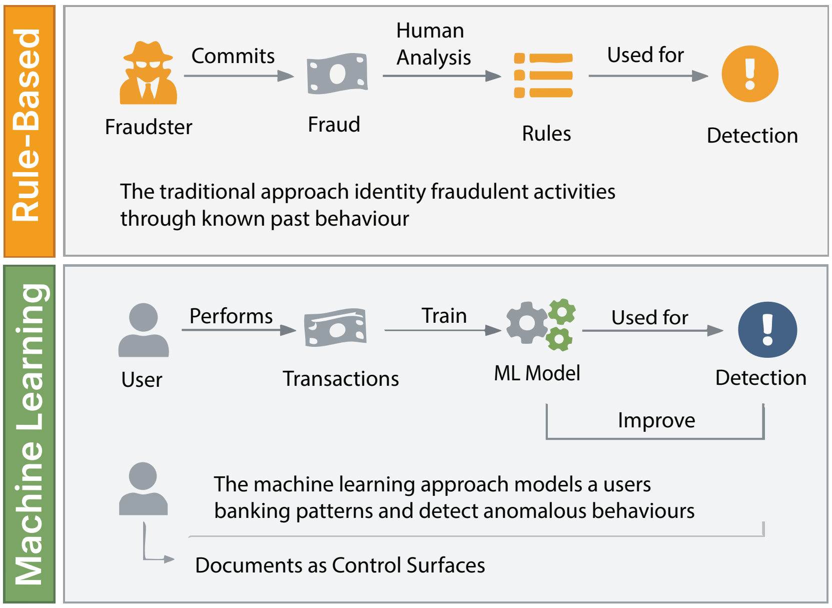

In agentic environments, memory becomes operational. Agents track evolving behaviour, update risk posture, and escalate based on accumulated context. Humans engage with richer narratives rather than raw alerts.

Fraud management shifts from chasing events to understanding behaviour. Rules still matter. Memory makes them relevant.

Why Agents Need Governance

As AI systems grow more capable, attention often shifts to intelligence. Models get better. Agents get faster. Outputs improve. Yet in **BFSI**, failures rarely stem from a lack of intelligence.

They stem from a lack of supervision. Supervision determines how systems behave under pressure. It defines when decisions pause, when they escalate, and who ultimately holds authority. Without it, even accurate systems can act at the wrong moment or with incomplete context.

This is why intelligence alone does not earn trust.

In supervised systems, judgment remains layered. Agents operate within bounds. Supervisor agents monitor outcomes across time. Humans intervene with visibility into both rationale and history. Decisions remain explainable long after they conclude.

In unsupervised systems, risk accumulates quietly. Errors propagate across workflows. Accountability blurs. When issues surface, reconstruction becomes difficult and confidence erodes.

Well-designed supervision reduces late-stage intervention by catching uncertainty early. It narrows escalation to where judgment is genuinely required. It ensures that learning feeds back into the system rather than staying with individuals.

For regulated environments, this design choice separates scalable intelligence from fragile automation.

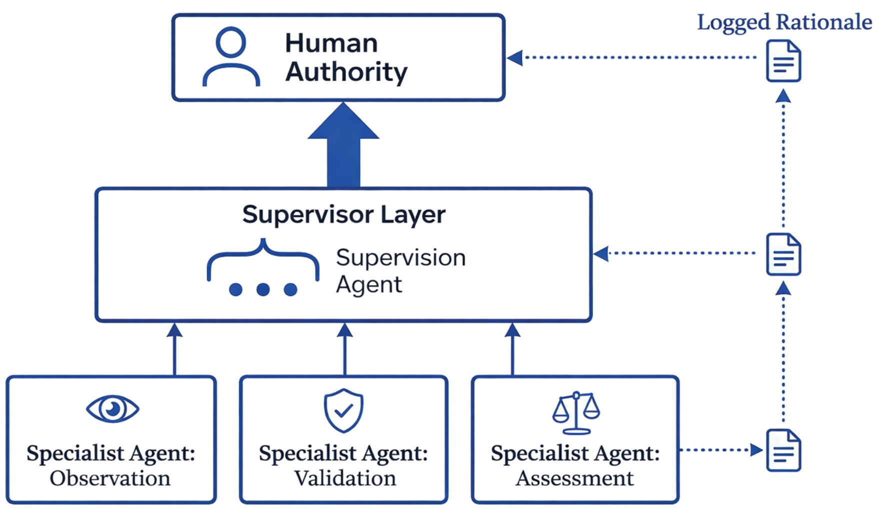

Control without micromanagement.

Inside an Agentic Architecture

The requirements are clear now.
Decisions must adapt over time. Context must persist across steps. Human authority must remain visible. What enables this is not a single model, but an architecture.

Agentic architectures organise intelligence around roles, not tasks. Each agent serves a specific purpose. One observes behaviour. Another validates information. A third assesses exposure. No one acts in isolation. They operate within a shared context and under supervision.

A system of Defined Roles

Each agent has a clear mandate. None acts in isolation.

Observation Agent:
Monitors incoming data and behavior.

validation Agent:
Verifies information against trusted sources

Assessment Agent:
Evalutes risk, exposure, or opportunity.

Supervision Agent:
Oversees agent interactions and flags anomalies

Roles interact through context, not sequence.

In an agentic system, decisions move through stages rather than steps. Signals accumulate. Confidence adjusts. Agents defer when uncertainty rises. Humans intervene with full context, not partial snapshots.

Agentic architectures allow **BFSI** organisations to scale judgment without losing control. Here's how it works in practise

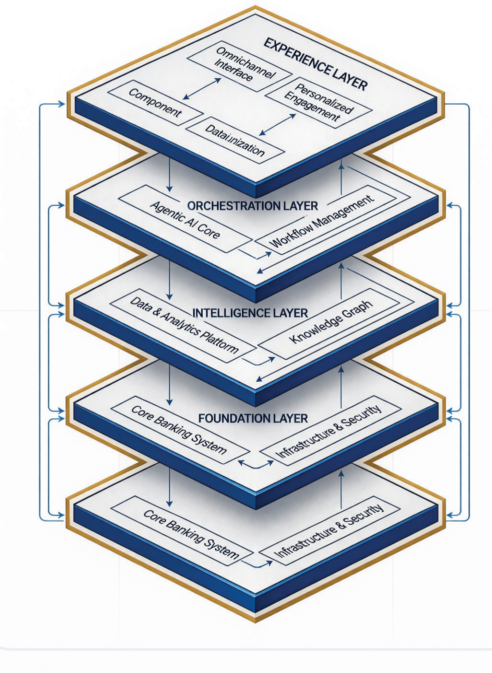

Layer 1 (Bottom): Intelligent Document Processing (GLIB)
Layer 2: Data Orchestration & Knowledge Graph
Layer 3: Agentic Orchestration
Layer 4: Business Outcomes

GLIB Agentic Architecture

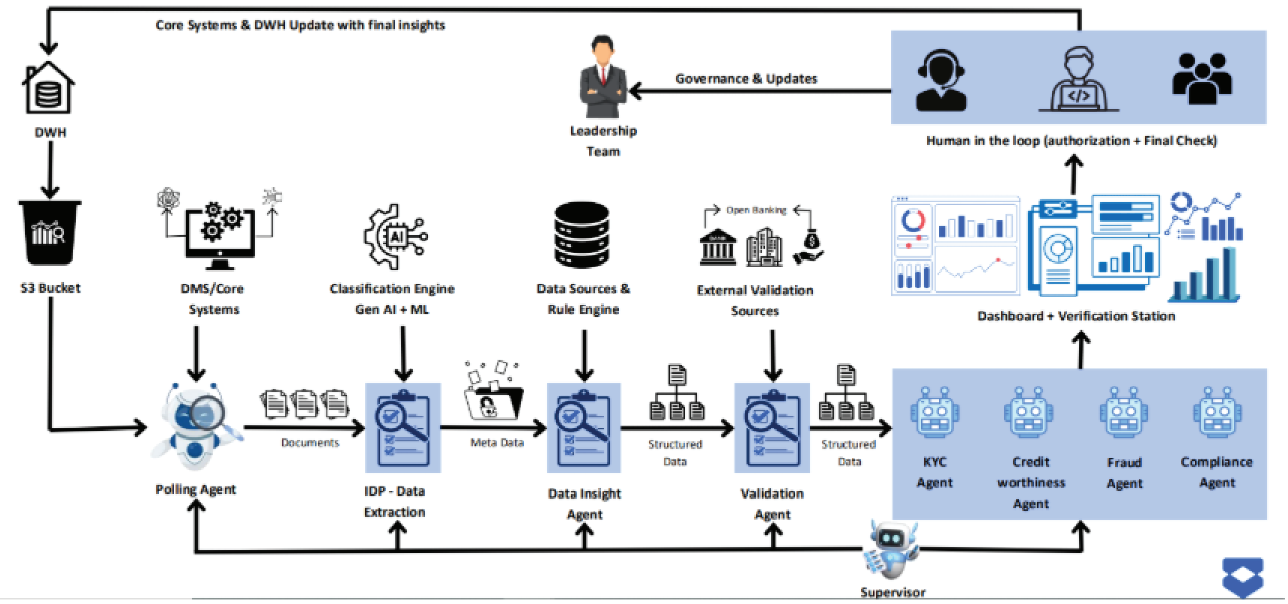

Why Regulation Actually Enables Agentic AI

The pace of GenAI adoption is increasing rapidly, with a 77% increase in organizations moving from pilot to scale. While the technology is clearly ready, the prevailing regulatory stance has been perceived as a hard stop, particularly for systems involving autonomous decision-making.

The paradigm is shifting. The risk tolerance is shifting from no autonomous decisions to supervised autonomous orchestration.

On 13 August 2025, the RBI released the **Framework for Responsible and Ethical Enablement of Artificial Intelligence (FREE-AI)** report. This framework is aimed at guiding innovation while managing risks associated with AI adoption in regulated financial institutions.

Banks and other financial institutions are encouraged to deploy AI with appropriate governance, controls, and oversight mechanisms :akin to how other regulated technologies (**e.g.**, credit scoring models, fintech APIs) are governed.
So, regulation enables supervised autonomy.

Multi-Agent Workflows in Credit

Credit, risk, and compliance each evaluate the same case through different lenses. Each view is valid. The challenge is alignment. In traditional workflows, convergence happens late, often through manual escalation. Context gets rebuilt. Timelines stretch. Outcomes vary. Multi-agent workflows change where convergence occurs.

Instead of waiting for humans to reconcile outputs, specialised agents collaborate through a shared context. One agent assesses cash-flow behaviour. Another evaluates exposure and concentration. A third validates compliance signals. Each updates the same decision state rather than producing disconnected outputs.

This design creates two advantages. First, convergence happens earlier. Conflicts surface while there is still time to adjust. Second, accountability remains clear. Each agent's contribution is visible. Rationale remains traceable. Humans step in to resolve ambiguity, not to assemble fragments.

Crucially, multi-agent workflows do not remove human judgment. They **focus** it. Analysts review narratives instead of raw data. Supervisors assess confidence instead of chasing completeness. Decisions move with fewer handoffs and clearer ownership.

Sample Multi-Agent Workflow

Credit Analyst Agentic AI

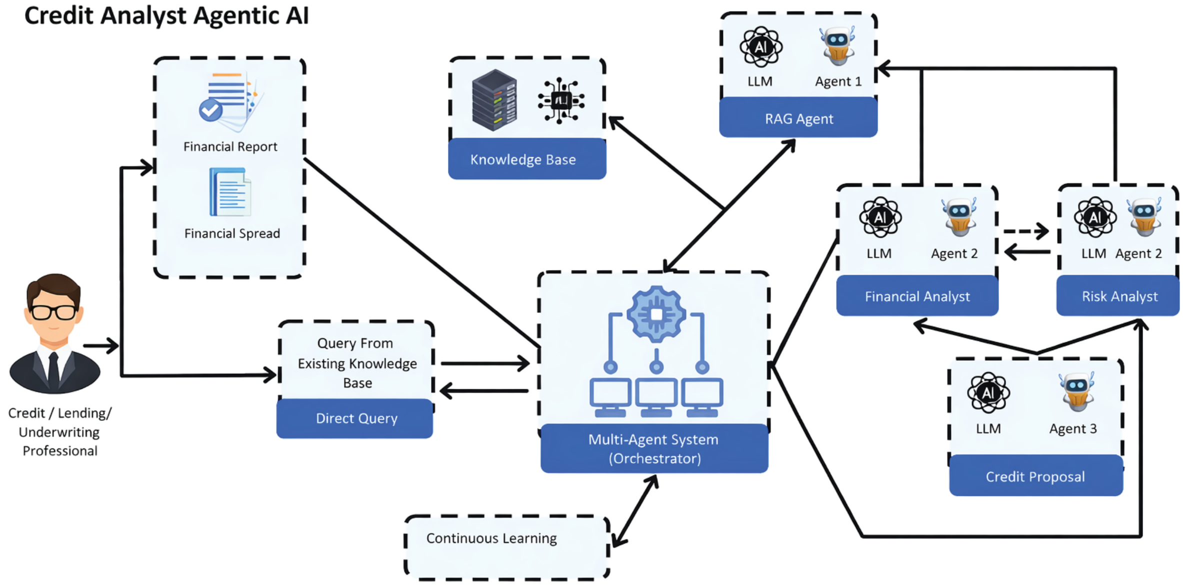

Lifecycle Thinking
Beats Use-Case Thinking

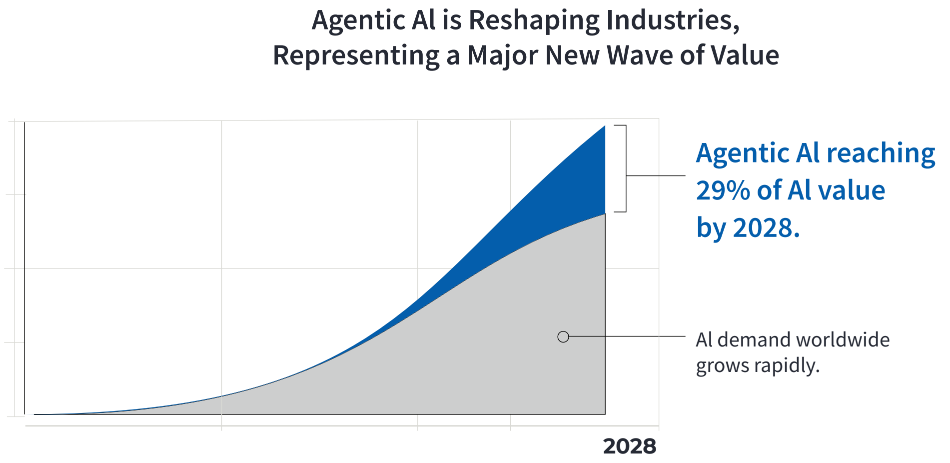

Use-Case Thinking

Goal: Solve an isolated task

Unit of Work: Discrete Project

Data & Learning: Siloed & Disposable

Value Creation: Linear & Additive

Outcome: A collection of tools

Mantra: "What can AI do for this problem?"

Lifecycle Thinking

Goal: Build a core capability

Unit of Work: Integrated System

Data & Learning: Shared & Persistent

Value Creation: Compounding & Exponential

Outcome: An intelligent ecosystem

Mantra: "How does this problem make our AI system smarter?"

Single Pilots Fail; Shared Memory Wins

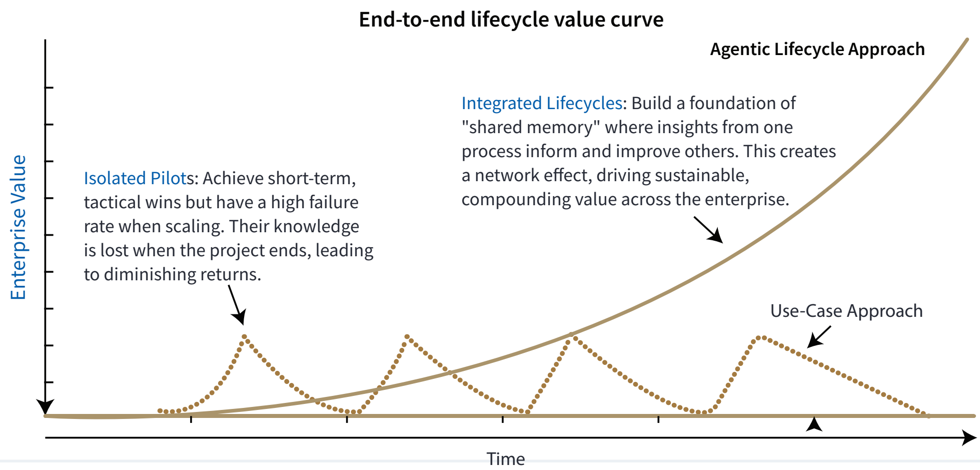

Agentic AI moves beyond simple automation. Coordinated intelligence through agentic AI creates a system of specialized AI agents that collaborate, share information, and reason together to execute complex, multi-step business processes.

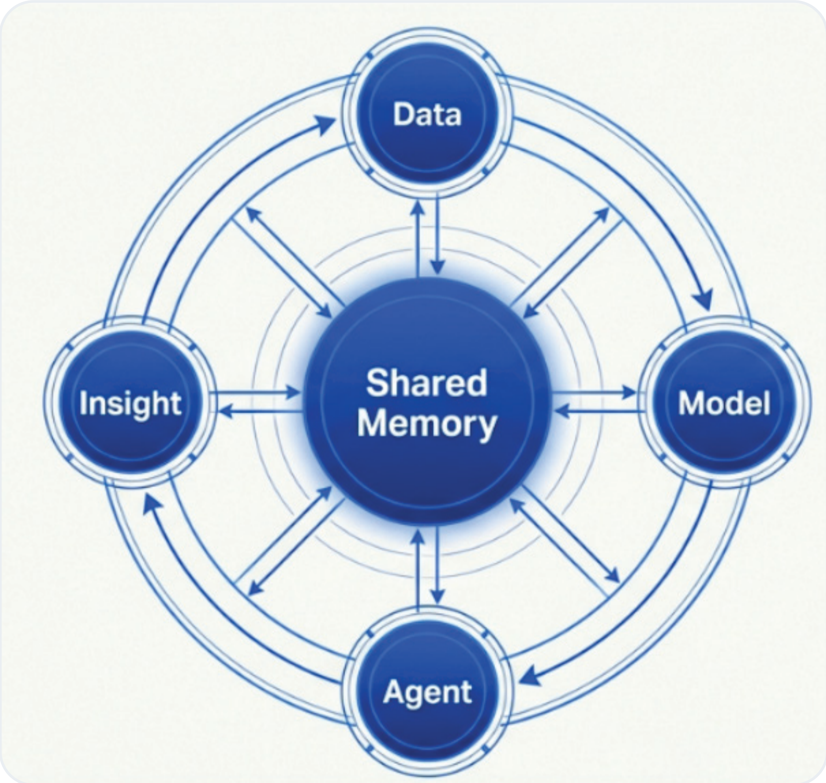

They are designed to work as coordinated teams instead of isolated tools. Together, they establish domain intelligence.

Agent 1 (Data Ingestion)
Gathers and structures applicant data from multiple sources.

→

Agent 2 (Risk Analysis)
Assesses creditworthiness using internal models and market data.

→

Agent 3 (Compliance Check)
Verifies the application against regulatory requirements.

→

Agent 4 (Decisioning)
Synthesizes inputs from other agents to recommend a final decision.

The Risk of AI Sprawl

Disparate AI project outcomes are common; **66% report tangible productivity gains**.*

But this masks a critical reality. A significant portion of the initiatives are failing to deliver meaningful **ROI**, and the gains that are realized are often tactical, not strategic.

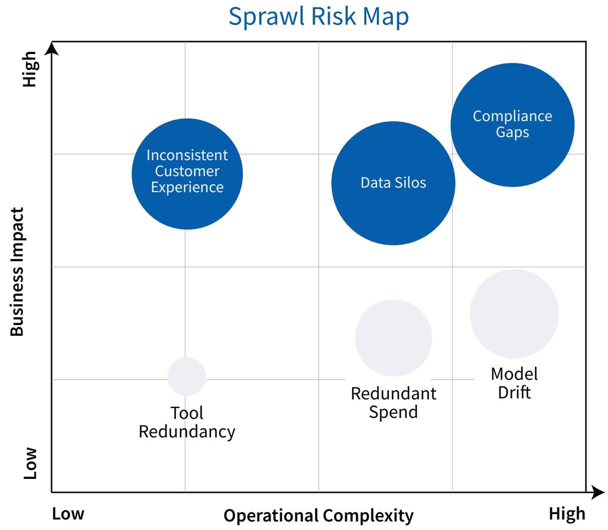

Uncoordinated AI is noise; organised agents deliver signal.

A Successful Roadmap Balances Three Core Variables

"Roadmaps must balance maturity, risk and impact."

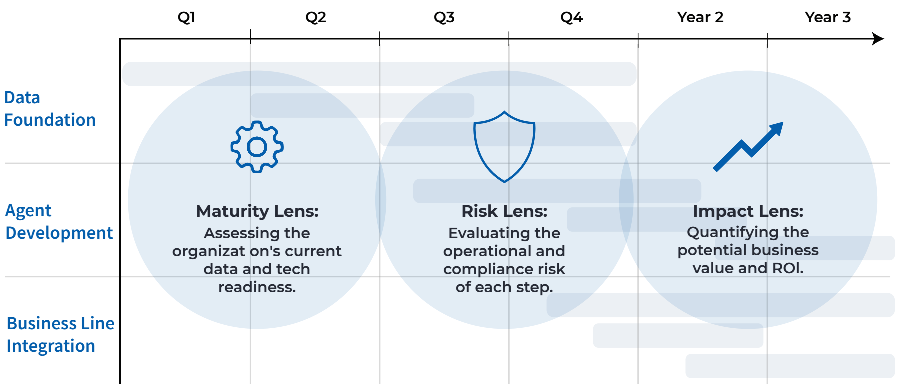

Example: **Tying the Roadmap to Market Strategy**

How can our roadmap position us to capture opportunity in high-growth markets, such as the projected AI in BFSI growth in India through 2033?

Building an Agentic Roadmap

Agentic capability is not built overnight through ad-hoc projects. It requires a strategic roadmap that sequences investments, develops foundational capabilities, and aligns AI development with long-term business objectives.

1. The Opportunity is Immense.

AI is fundamentally reshaping banking. Realizing its full value is not guaranteed and requires moving beyond isolated experiments.

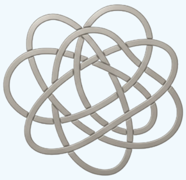

2. The Risk of Sprawl is Systemic

Uncoordinated, use-case-driven AI is the single biggest threat to your AI ROI, creating noise, risk, and inefficiency.

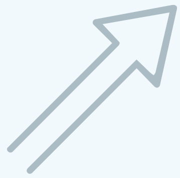

3. The Path Forward is an Agentic Roadmap.

A lifecycle-driven strategy, built on multi-agent collaboration and shared intelligence, is the key to unlocking sustainable, compounding value.

This is the bridge from ambition to execution.

Why Indian BFSI Localisation Matters

USD 8,090 M

Projected market size for Artificial Intelligence in the Indian **BFSI** sector by 2033.

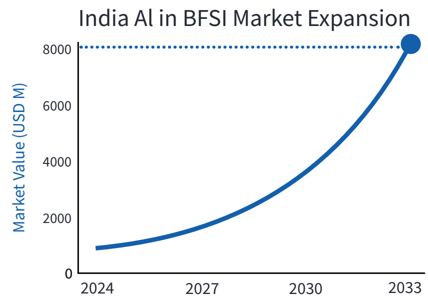

The AI opportunity in Indian banking is accelerating.

Success with agentic AI demands a new model built on two pillars:
- **Hyperlocalization**: Designing solutions for India's specific complexities
- **Outcome-based scale**: Measuring success in tangible business results.

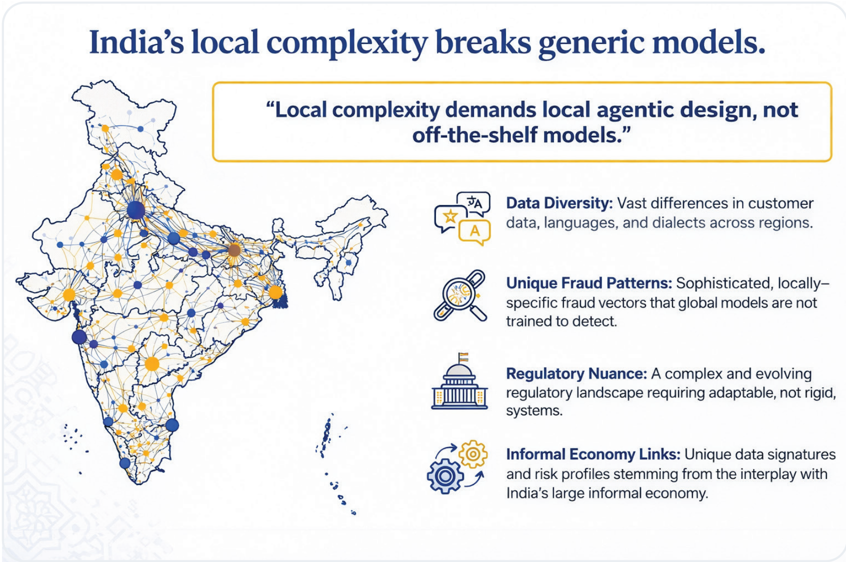

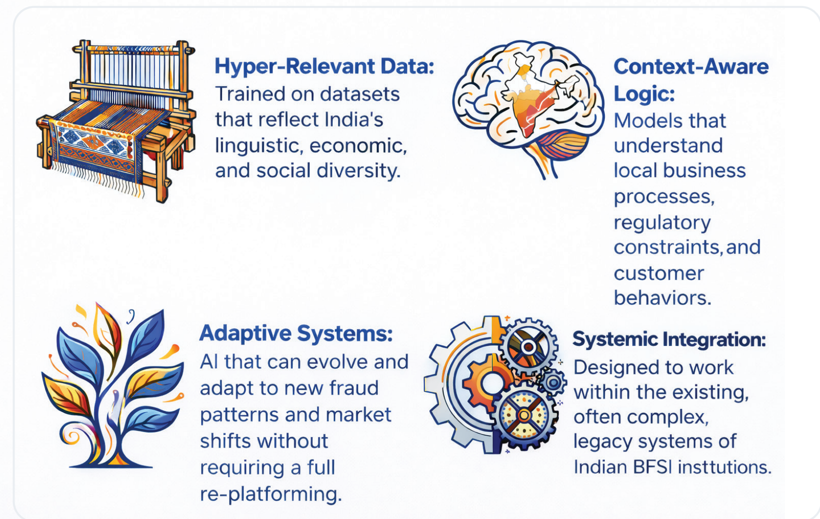

Indian financial institutions need

This is exactly what Glib Agentic AI enable:

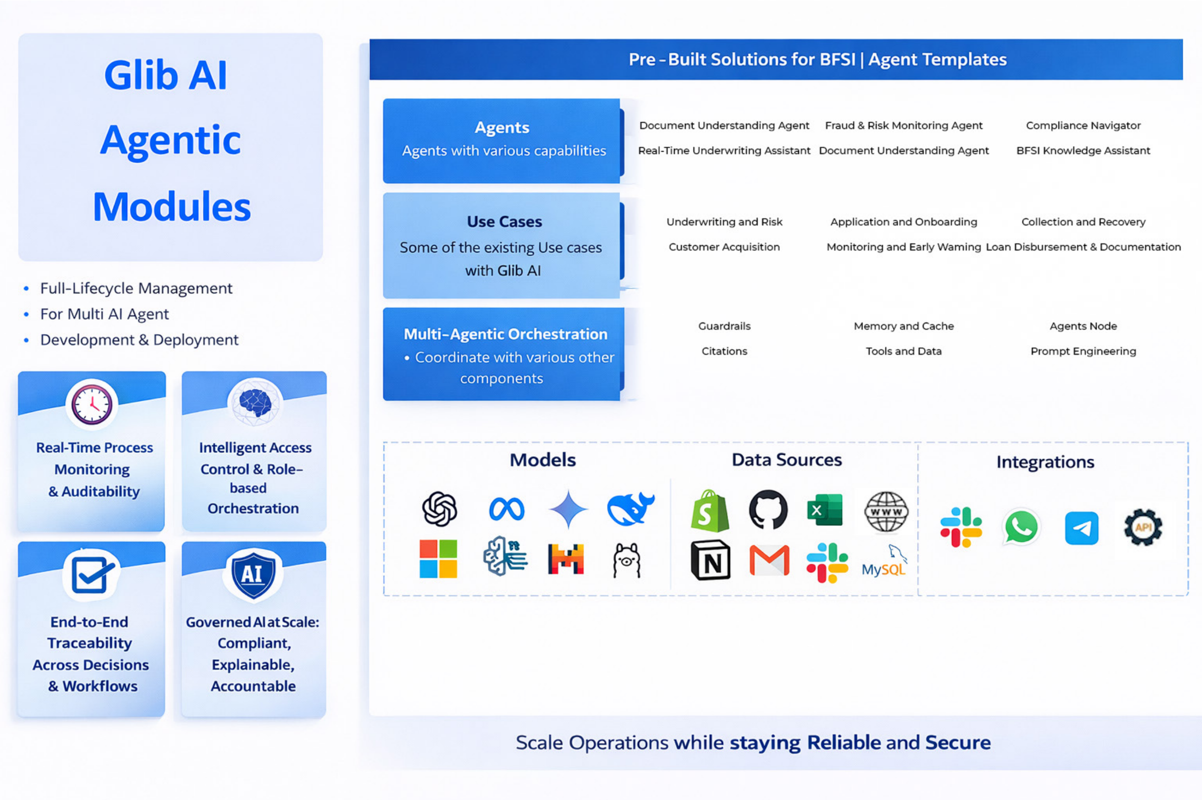

Case Signals:
What Actually Scales

Industry AI signals of fraud & efficiency improvements validate systemic change.

The Old Metric of Scale (Activity)	The True Metric of Scale (Outcomes)
• Number of licenses purchased	✓ Percentage reduction in fraud incidents
• Models deployed in production	✓ Decrease in loan processing time (hours)
• Departments using the new tool	✓ Increase in customer query resolution rate
Result: 'Activity' is high, but business impact is ambiguous.	**Result: The AI's contribution to P&L and operational efficiency is direct and measurable.**

Early adopters who have shifted from generic tools to localized, outcome-driven AI are seeing significant, measurable improvements. The evidence is not in marketing claims, but in the transformation of their operational dashboards: from lagging indicators of activity to real-time indicators of value.

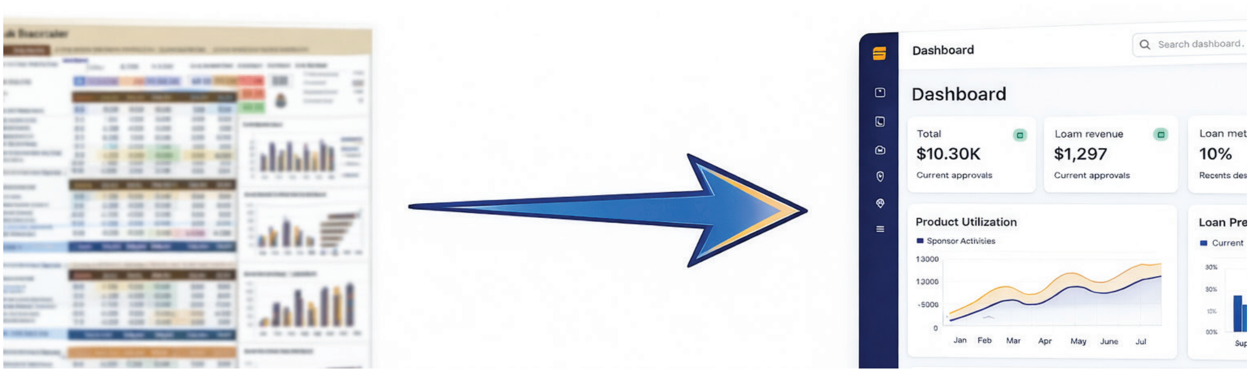

Note: A short GIF showing how Glib dashboards are built can be shown here

Case Evidence: Large African Bank: Lending Automation

The Constraint: **ABSA Bank** processed 75,000+ loan applications annually through manual workflows. Each application required 4.5 hours (document review, MCIB credit bureau fetch, credit scoring). TAT: 6-8 hours for simple cases, 2-3 days for complex profiles. Annual burden: 337, 500 FTE-hours (~162 analysts).

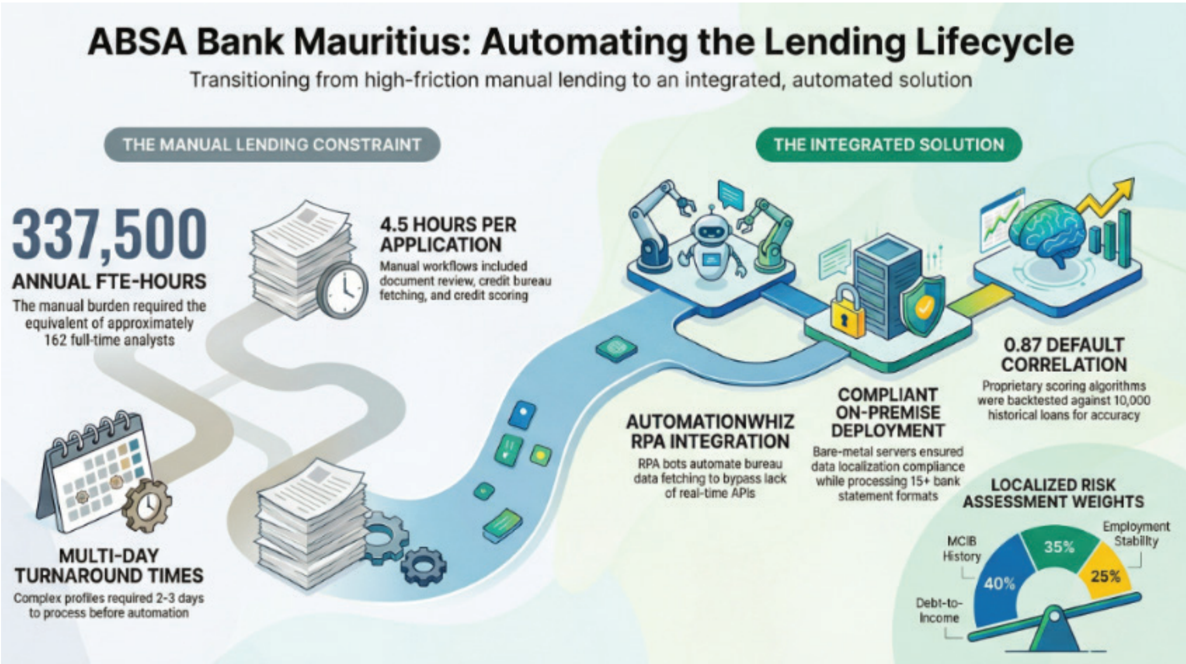

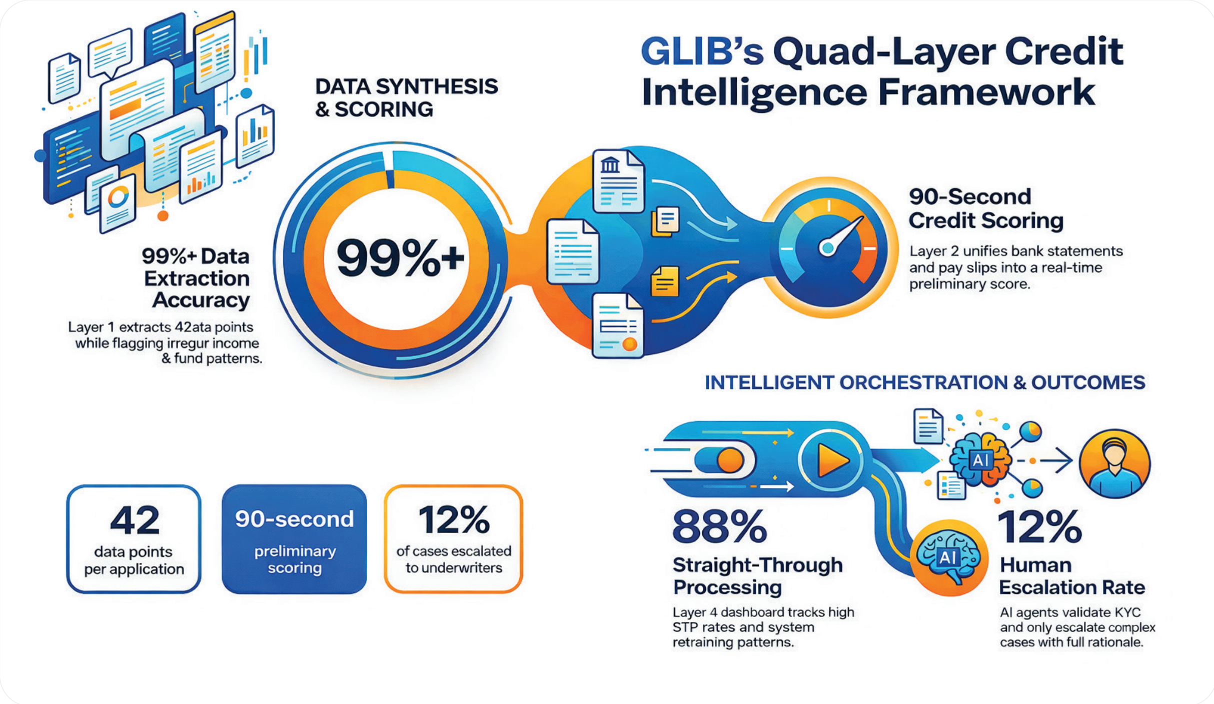

Results:

- **TAT**: 4.5 hours → 2 minutes (98.5% reduction)
- **Annual hours saved**: 318,750 (~153 FTEs redeployed)
- **Manual effort**: 80% reduction
- **Business unlock**: Quick lending product launched—+18% YoY unsecured lending portfolio
- **Implementation**: 6-month phased rollout. 40 underwriters retrained. Bank of Mauritius regulatory audit passed. Quarterly retraining from human override feedback.

Case Evidence: India's Largest Life Insurance Company: Document Processing at Scale

The Volume Challenge: SBI Life processes 50+ million pages annually (death claims, maturity claims, policy servicing). Phygital chaos: Regional operators scan 10-15 docs into single 40-page PDFs. Death certificates from rural hospitals (handwritten, faded). Bank statements (200+ formats). KYC docs (Aadhaar, PAN—photocopied multiple times).

Baseline: 100 FTEs manually classified, extracted, validated. TAT: 4-6 hours per claim (8-12 million FTE-hours annually). Peak volumes (post-pandemic) saw 48+ hour delays. IRDAI scrutinized.

Phase	Challange	What Broke	How GLIB Fixed It	Outcome
Phase 1: Pilot (Months 1-3)	Death cert chaos	40+ variations, handwritten regional languages. Initial OCR: 65% accuracy	Collected 5K rural samples, retrained with fuzzy date matching	92% accuracy
	Bank statement explosion	Expected 50 formats, found 187 (legacy PSU redesigns, cooperative bank handwritten entries)	Shifted to adaptive OCR (ML learns format patterns dynamically)	94% accuracy
	Cross-validation false positives	"Rajesh Kumar" vs. "Rajesh K." flagged 30% unnecessarily	Fuzzy name matching (Levenshtein distance, phonetic matching)	8% false positives
Phase 2: Scale (Months 4-9	Peak load crashes	Diwali volume spike (250K claims/month) caused 72-hour backlogs	Load balancing across 3 data centers, batch prioritization	12-hour max backlog
	Aadhaar masking issue	Post-2022 UIDAI mandates broke validation (masked vs. full number mismatch)	Multi-factor matching (name + DOB + last 4 digits + account). IRDAI approved	100% compliance
	Human resistance	Veterans bypassed HILT—60% adoption only	40 training sessions, gamified accuracy scores, showcased STP success	92% adoption
Phase 3: Hardening	Edge case coverage	200+ real-world variations discovered in production	Urdu script OCR, zero-transaction logic, pencil-form preprocessing,	Production ready at scale
(Months 10-18)			joint account matching, post-dated cheque grace periods	Death cert chaos

Points scored

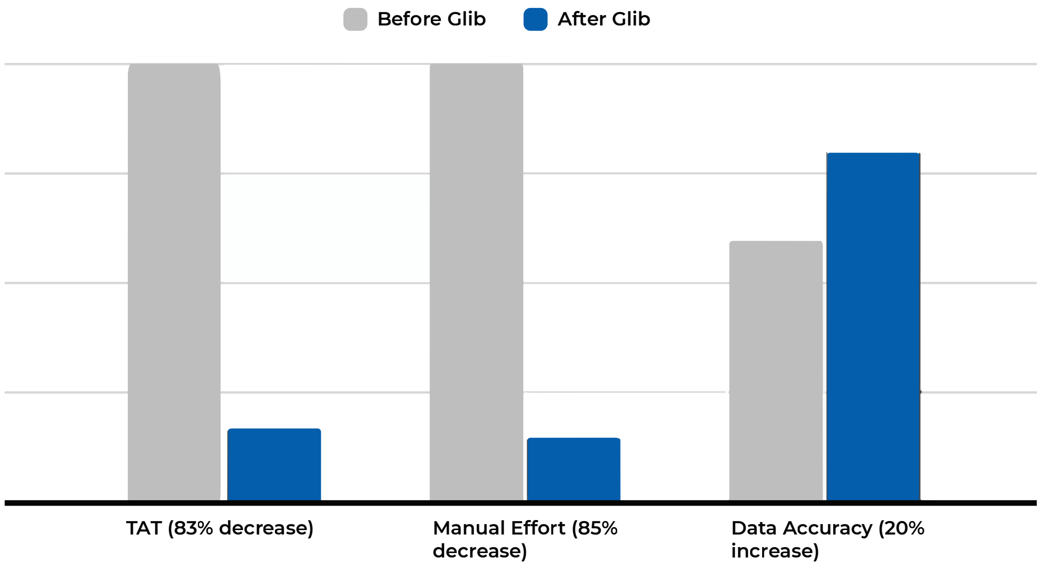

Business Impact:

TAT: 55 sec median (STP), complex cases 45 mins (vs. 4-6 hours)
FTE: 100 → 15 (85% reduction); volume scaled 50M → 65M pages (+30%) with no headcount increase
Accuracy: 80%+ end-to-end; system absorbed peak loads without manual intervention

Case Evidence: Global Ratings Analytics Organisation : Financial Data Automation

The Analyst Trap: CARE Ratings' 50 analysts spent 4-6 hours per issuer—2-3 hours manually typing financial data from PDFs into Excel, 2-3 hours calculating 800+ CAM parameters (leverage, liquidity, profitability ratios). Max capacity: 24,000 issuers/year. Analysts' time: 80% data entry, 20% credit judgment. Attrition: 25%/year ("I didn't join to copy-paste").

Workflow Transformation:

Before (6-8 hours):

1. **Manual extraction (2-3 hrs)**: Balance sheet, P&L, cash flow, notes from 80-150 page PDFs
2. **Chart of accounts mapping (1 hr)**: 60+ items to CARE's proprietary codes
3. **CAM prep (2-3 hrs)**: Calculate 800+ parameters manually, write narrative
4. **Senior review (1 hr)**: Spot-check 10 items, flag discrepancies, send back (30% rework rate)

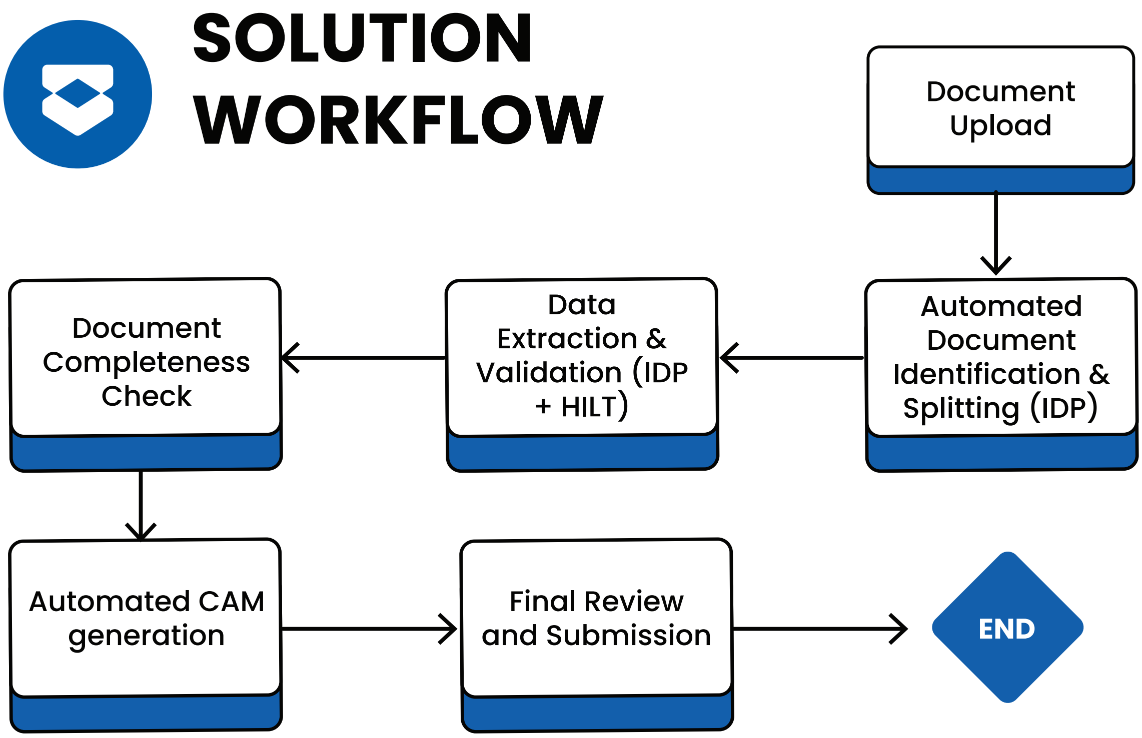

After with FinRay (30 mins):

1. **Auto-upload + extraction (5 mins)**: AI-powered OCR extracts from 40+ Indian GAAP/Ind-AS formats, handles merged cells/footnotes. Chart-of-accounts auto-mapping (97% accuracy).

2. **Validation (2 mins)**: Balance sheet reconciliation, anomaly detection (15 red flags—debt spikes, negative CFO, related-party transactions >10% revenue, auditor qualifications).

3. **AI CAM generation (3 mins)**: 800+ parameters calculated instantly. Preliminary CAM with 5-year trends, sector benchmarks, risk highlights auto-generated.

4. **Analyst review (15 mins)**: Senior validates AI flags, adds qualitative judgment (management quality, industry outlook), adjusts ratings.

5. **Approval (5 mins)**: Rating committee focuses on strategic discussion (not data verification).

Results:

- **Analysts**: 50 → 12 (76% reduction, ₹26 crore annual savings)
- **TAT:** 4-6 hours → 30 mins (8-12× faster)
- **Capacity:** 24K → 36K issuers/year (+50%, no FTE growth)
- **Accuracy:** 88% → 99%+ (65% improvement)
- **Market share:** 28% → 32.2% (+4.2pp vs. CRISIL/ICRA)
- **New revenue:** ₹18 crore (SME segment unlocked)
- **Implementation:** 12 months.
- **Edge cases:** Consolidated financials (15 subsidiaries), foreign currency translation, aggressive accounting detection. SEBI audit passed.

From Questions to Decisions for CXOs

Market momentum and adoption trends point to systemic change.

The strategic shift is unmistakable: The advantage moves from speed to system-level intelligence. Agentic AI supports building a system where autonomous decisions are governed, accountable, and aligned with enterprise goals.

The five questions posed earlier in this report—authority, data integrity, workforce transformation, capital efficiency, and competitive timing—are not independent technical hurdles. They converge into a single design problem: how to structure intelligence so that autonomy and accountability scale together.

The organisations making progress have stopped treating agentic AI as a procurement decision. Liability concerns dissolve when systems preserve traceability by design. Data quality stabilises when extraction precedes reasoning, not the other way around. Workforce anxiety diminishes when automation targets drudgery, not expertise. ROI becomes measurable when outcomes, not activity, define success. Competitive risk shifts when regulation enables governance rather than blocking autonomy.

RBI's FREE-AI framework reframed the conversation. Supervised autonomy is not just permissible—it is expected. The question for Indian BFSI is no longer whether agentic systems belong in decision workflows. It is whether institutions will design them deliberately or inherit them reactively.

The deployments detailed in this report show what deliberate design delivers.

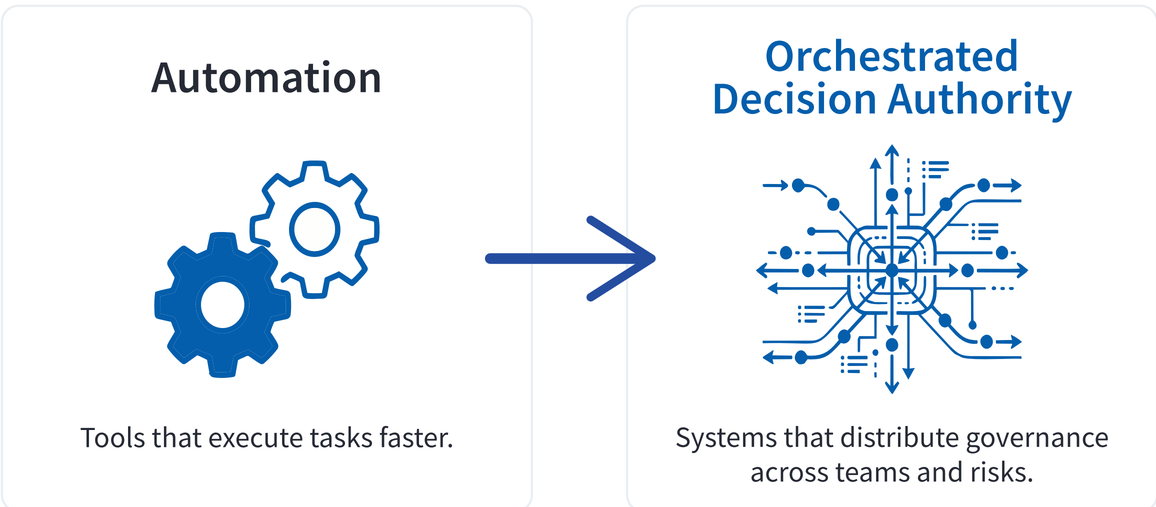

GLIB's 4-layer architecture enabled a bank processing 75,000 applications annually to compress decisioning from hours to minutes while maintaining full audit trails.

The same framework automated 50 million pages under regulatory scrutiny for an insurer without sacrificing accuracy. A ratings agency using GLIB's orchestration layer reassigned analysts from data entry to interpretation, compressing turnaround time by 60% in year one.

None began with technology. All began with clarity about what decisions matter, where judgment concentrates risk, and how accountability must flow when systems operate autonomously:—the foundation-to-outcome structure GLIB's framework makes explicit. .

Glib applies a proven framework to your operational challenges. We employ a systematic approach to identify and automate critical process bottlenecks. **Connect with us today**!

Source Links

- https://www.thewallstreetschool.com/blog/ai-in-india-financial-firms-2025/

- https://www.thewallstreetschool.com/blog/ai-in-india-financial-firms-2025/

- https://www.scielo.br/j/pope/a/VKZKqQxPVx9wS7VHWHVLxCt/

- https://www.ey.com/en_us/insights/banking-capital-markets/ai-in-banking-ey-parthenon-genai-survey-insights

- https://www.financialexpress.com/business/banking-finance-74-of-indian-financial-firms-adopt-genai-ey-report-3775455/

- https://www.ibef.org/news/generative-ai-set-to-improve-banking-operations-in-india-by-46-rbi-report

- https://www.precedenceresearch.com/artificial-intelligence-in-banking-market

- https://www.bcg.com/publications/2025/branches-to-bots-will-ai-transform-retail-banking

- https://economictimes.indiatimes.com/industry/banking/finance/banking/ai-could-reshape-half-of-the-roles-in-indias-banking-sector-report/articleshow/123504418.cms?from=mdr

- https://invest.up.gov.in/wp-content/uploads/2024/05/AI-adoption_130524.pdf

- https://www.imarcgroup.com/india-artificial-intelligence-in-bfsi-market

- https://indiaai.gov.in/article/rbi-study-explores-how-indian-banks-leverage-ai

- https://www.business-standard.com/industry/banking/almost-three-fourths-of-banks-nbfcs-implement-chatbots-rbi-s-patra-124012900876_1.html

- https://www.caspianone.com/ai-in-financial-services-report

- https://www.mckinsey.com/industries/financial-services/our-insights/capturing-the-full-value-of-generative-ai-in-banking

- https://www.pwc.com/us/en/tech-effect/ai-analytics/ai-agent-survey.html

- https://www.imarcgroup.com/india-artificial-intelligence-in-bfsi-market
1