

Unlocking Enterprise Intelligence:

A Journey of Innovation and
Transformation with Agentic AI



Foreword by Manish Jain



Artificial Intelligence is moving from a tool that predicts and assists to a system that can increasingly interpret intent, coordinate tasks and act within defined boundaries. In India's financial services ecosystem, the impact is profound. Over the past decade, digital infrastructure expansion, access to formal credit and innovation across banks, NBFCs and fintechs have transformed how consumers discover, access and manage financial products. Artificial Intelligence now has the potential to reshape the next part of that journey not simply by making existing processes faster, but by changing how consumers and institutions interact, make decisions and act.

Consumers are already beginning to use AI to understand financial concepts, compare alternatives, manage their finances and assist in making decisions on their behalf. At the same time, they are concerned about inaccurate information, bias towards specific providers, misuse of personal data and the possibility of manipulation. This combination of enthusiasm and caution suggests that the next phase of adoption of AI in financial services will be defined not by how autonomous technology can become, but by how confidently consumers and financial institutions can trust it.

Agentic AI takes this discussion further. Agentic systems can potentially interpret an objective, bring together relevant information, coordinate multiple activities and initiate permitted actions. In financial services, this could fundamentally change the journey across all the stages of the credit life cycle, from acquisition to overall portfolio management, to ensure a seamless customer experience.

But intelligence alone is not enough. Financial decisions depend on reliable data, robust risk assessment, clear policy and accountability. **For lenders, this creates an opportunity to rethink how intelligence moves through the organisation.** Today, data, models, decision strategies and workflows can remain fragmented across different systems and teams. However, orchestration through connected and governed platforms has the potential to bring these elements together - helping financial institutions move more efficiently from dynamic risk signal to analysis, validation, approved decision-making capabilities and continuous monitoring.

This report looks beyond AI as a technology trend. It explores what it means for consumers, lenders and the broader credit ecosystem as financial services become increasingly agent enabled.

That is the opportunity ahead: to build an agent-ready financial ecosystem in which technology helps both consumers and financial organisations make better decisions, while trust remains the foundation on which those decisions are made.

Manish Jain

Country Managing Director
Experian India



Executive Summary

From digital lending to the agent-ready credit enterprise

India's lending market is entering a new phase shaped by the convergence of digital ecosystems, expanding data availability, advanced analytics and rapid progress in artificial intelligence. The shift is moving beyond isolated automation towards connected decisioning: intelligent systems that can interpret context, coordinate approved workflows, act within defined boundaries and return outcomes to monitoring and learning. The opportunity is not maximum autonomy, but governed orchestration — using AI to accelerate the journey while lenders retain authority over affordability, fraud controls, policy, pricing and the final credit decision.

The consumer side of this transformation is already visible. In Experian's July 2026 research with Forrester Consulting among 480 credit-active, digitally literate Indian consumers,

- 89% trust LLMs to compare loans across different providers while 60% are comfortable with an AI agent applying for a loan or credit card on their behalf.
- Yet trust is decisive — 85% prefer an LLM provided by or connected to a financial provider they already trust.

Key Takeaways

- Consumers are not only using AI tools to make informed financial decisions; they also show a clear preference for tools connected to trusted providers and designed specifically for financial services.
- As AI tools make product comparison easier, faster and more personalised, financial services providers may face a more competitive market.
- There is an opportunity to provide trusted AI that combines intelligence with reliable information and appropriate financial context.
- The more an agent can do, the stronger the trust infrastructure around it needs to become.
- Agentic technology, delivered through connected platforms, could reduce several forms of friction simultaneously.
- The opportunity is to leverage platforms to automate the model lifecycle, while keeping accountability with responsible human teams.
- As financial services become more agent-enabled, trusted connectivity between AI agents and institutions could become as important as the intelligence of the agents themselves.

The institutions that create lasting advantage will not be those with the most AI pilots or the greatest degree of autonomy. They will be those that become truly agent-ready: discoverable and responsive, connected across data and decisions, and disciplined about evidence and accountability.



Index

01	Key Insights At A Glance:	5
02	Chapter 1: The Rise of the Agentic Consumer – A Study of Transforming Financial Behaviour	11
03	Chapter 2: The Role of Agentic AI in the Indian Credit Landscape	19
04	Chapter 3: The Personal Loan Story – Focus on Fintech-led Growth	26
05	Chapter 4: Trust at Scale – Data Governance and Regulatory Checks for Agentic AI	31
06	Chapter 5: Harnessing the Power of AI through Platforms	36
07	Chapter 6: The Road Ahead	43



Key Insights At A Glance

1. Consumers are not only using AI tools to make informed financial decisions; they also show a clear preference for tools connected to trusted providers and designed specifically for financial services.


Indian consumers are quickly adopting AI for practical financial decisions. Large Language Models (LLMs) are being used to understand financial terms, compare offers and manage personal finances – helping consumers save time, reduce effort and make better choices. As per Experian's report in partnership with Forrester (where 480 credit-active Indian consumers were part of a survey), various opinions from Indian consumers regarding usage of AI agents have come to light.

But as AI becomes more embedded into financial decision-making, trust emerges as the true differentiator.

For financial institutions, this creates both responsibility and significant opportunity. The opportunity is not simply to deploy more AI, but to offer trusted, contextual intelligence that combines the conversational ease of the generative AI with reliable real-time product information.

The next competitive question is not simply whether a provider offers AI, but whether customers can trust the intelligence behind it.

Snapshot of key findings




trust LLMs to compare loans across different providers



are comfortable with an AI agent applying for a loan or credit card on their behalf



would feel more comfortable using an LLM that is provided by or connected to a financial provider they already trust



would be comfortable giving an AI agent conditional or full autonomy for a loan or credit card application



would feel more confident in their purchase decisions knowing that AI has searched for the best offers

Base: 480 credit-active, digitally literate Indian consumers

Source: Experian research conducted by Forrester Consulting, July 2026




2. As AI tools make product comparison easier, faster and more personalised, financial services providers may face a more competitive market.

AI is removing friction from comparing credit products. This could shift the battleground for competition. The first interaction may happen through a consumer-controlled AI agent rather than on a lender or marketplace website. To be considered in the moment, providers must be discoverable with current product and eligibility information, comparable in clear terms and responsive with fast, reliable pre-qualification.

AI-driven competition will not make price the only differentiator, it will make the complete proposition – rates, fees, eligibility, etc. more transparent and easier to evaluate. Providers that can consistently deliver accurate information, respond quickly and offer the right product for the consumer's needs, will be better positioned to win.

The competitive question is shifting from “How do we bring the customer to our channel?” to “How do we ensure our proposition is discovered, understood and considered when the customer's AI agent is searching the market?”



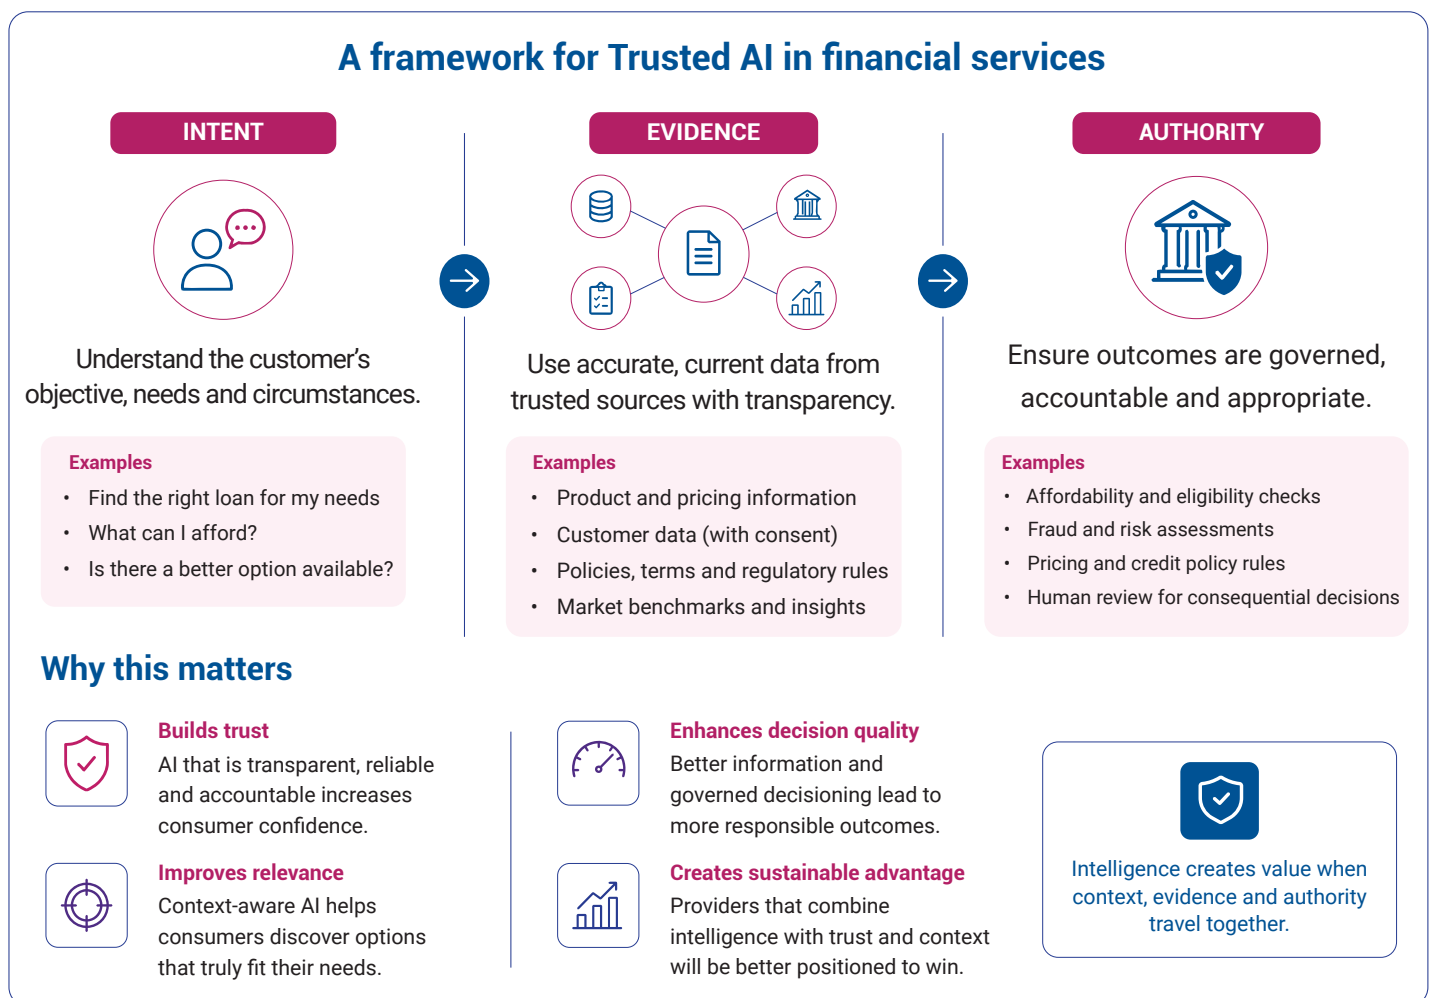
3. There is an opportunity to provide Trusted AI that combines intelligence with reliable information and appropriate financial context.

Consumer's financial decisions require more than answers – they require context, accuracy and accountability. A consumer may start with a simple objective – “What can I responsibly afford?” An AI agent can help interpret that intent, but delivering a responsible outcome requires verified product information, permitted data, eligibility criteria and approved risk and policy frameworks.

As AI moves from interpretation to recommendation and action, the need for trusted financial context becomes even more important. For institutions, this creates an opportunity to build AI experiences differentiated by trust. Provider-connected AI can combine conversational simplicity with authenticated product information, appropriate customer context and existing controls.

Trusted AI requires three elements to work together: Intent, Evidence, Authority.

Trusted AI is created when useful intelligence is anchored to reliable evidence and clear decision authority.




4. The more an agent can do, the stronger the trust infrastructure around it needs to become.

As AI agents move from helping consumers with information on financial products towards initiating financial actions on their behalf, the nature of risk changes significantly. Greater delegation introduces new questions around what information the agent can access, which actions it is authorised to take and who remains accountable when something goes wrong. **This means that trust cannot be treated as a single control applied equally to every AI interaction.** The safeguards surrounding an agent need to increase in proportion to the authority and potential consequence of the action.

For financial institutions, the objective should therefore not be maximum autonomy but **appropriate autonomy supported by proportionate controls**. An agent should only progress from recommending to acting when the institution can demonstrate that the necessary consent, security, evidence, decision controls and customer recourse are in place.

In essence, agent capability and trust infrastructure need to advance together. Capability should expand only when consent, controls, evidence and human recourse expand with it.




5. Agentic technology, delivered through connected platforms, could reduce several forms of friction simultaneously.

Agentic technology has the potential to create value not simply by automating individual activities, but by connecting multiple activities across a financial journey. Consumers today may have to move through a long and disconnected process, often repeating information or having to navigate through different systems at each stage. An AI agent operating through a connected platform could potentially co-ordinate these activities, preserve relevant context and help move the customer more efficiently towards the end stage.

The value of agentic AI is multi-dimensional: it can potentially reduce the time and effort required to search and progress ahead in a particular process, the complexity involved in evaluating choices and the risk of overlooking important information. For consumers, this could create a fundamentally different financial experience.

The opportunity is equally significant for financial institutions. Today, data acquisition, verification, underwriting, servicing and portfolio management can operate as separate processes. Agentic AI can potentially co-ordinate approved tasks across these stages, while a connected platform provides the base required to support them.

The value of Agentic AI is not just faster communication. With connected platforms providing the analytical foundation, institutions can connect data, decisions and actions to deliver contextual progress from intent to an approved next step.



6. The platform opportunity is to automate the mechanics of the model lifecycle – not the accountability of model judgement.

As credit portfolios become more dynamic, financial institutions need to respond faster to growth and emerging risk signals. Yet the process of converting an insight into a production-ready model or strategy can remain sequential and intensively dependent on manual interventions. Each activity may be individually robust, but repeated movement between teams, tools and environments can slow the journey and create unnecessary rework, version risk and loss of context.

The important distinction, however, is between automating workflow mechanics and automating judgement. Agentic AI can prepare information, run approved analytical tasks, organise evidence, generate documentation and coordinate workflow steps. It should not independently redefine the business objective, choose policy without approval, waive validation requirements or authorise material production changes. Those responsibilities remain with accountable business, risk, model-validation and governance teams.

Experian Ascend provides the opportunity to connect these activities through a common intelligence environment. Data, analytics, models, artefacts and monitoring can remain connected and provide a seamless deployment option, all within the same platform.

The next era of lending will be shaped by agent-ready institutions that respond quickly and transparently while keeping data, decisions and accountability governed.

7. As financial services become more agent-enabled, trusted connectivity between AI agents and institutions could become as important as the intelligence of the agents themselves.

The emergence of the agentic consumer creates a new requirement for financial services. Consumers may increasingly use AI agents to research products, compare lenders, organise information and potentially initiate parts of a financial journey. But for these interactions to progress beyond information and recommendation, agents need a secure and governed way to connect with the data, services and systems of financial institutions.

Model Context Protocol (MCP) provides a common framework through which AI agents can connect with enterprise data, analytical services, business applications and decisioning systems. Rather than creating a separate integration for every agent and system, a common protocol can help establish how permitted information and approved capabilities are made available to an agent.

Agent-led finance will depend not only on smarter agents, but on trusted protocols that allow agents and financial institutions to interact safely, transparently and within clearly defined boundaries.



Chapter 1



The Rise of the Agentic Consumer –
A Study of Transforming Financial Behaviour

India's financial services landscape is entering a new phase of digital transformation. Consumers have become accustomed to discovering products online, completing applications digitally and managing financial relationships through mobile interfaces. Artificial Intelligence (AI) is now beginning to change not only how consumers access financial services, but how they make financial decisions helping them interpret information, compare alternatives and navigate increasingly complex choices.

Agentic AI could take this evolution further.

Rather than simply responding to individual questions, AI agents can potentially understand a consumer's objective, search and compare options, organise information, recommend actions and, within defined permissions, complete parts of a financial journey. This creates the emergence of the **Agentic Consumer**: a customer who remains in control of the financial objective and increasingly uses intelligent systems to move from intent to informed action.

A Deep Dive into the Indian Consumer's Perspective

In an Experian research study conducted in partnership with Forrester Consulting in July 2026, 480 credit-active, digitally literate consumers in India were asked for their views on the extent to which they are using LLMs like ChatGPT, Gemini and Claude, as well as how comfortable they are delegating a range of tasks to AI agents.

AI agents for the respondents are designed as general-purpose LLMs that can independently perceive information, reason through problems by understanding context and then take action across multiple steps. Although this type of agentic AI is still in the early stages of adoption, Experian set out to better understand how ready consumers are for agents and, importantly, what is needed to build trust and acceptance in using them.


A. Consumers Remain Optimistic Despite Financial Pressure

The starting point is a consumer who is broadly confident about their financial position, but who is also navigating meaningful financial pressures. The research shows that most consumers are feeling positive about their financial position and their ability to manage existing debt. However, there are clear signs of financial pressure, with many households juggling multiple credit commitments and relying on credit to cover essential expenses.

Around a quarter of respondents are showing signs of financial distress – this points to a clear need for more proactive, personalised support to help this segment stay in control.



Overall consumer financial health is positive



Credit appetite for the year ahead

Two-thirds (66%) of consumers would consider taking out additional credit in the year ahead, although current interest rates are acting as a deterrent. As a result, just over two-thirds (68%) would take on new credit only to meet essential needs, rather than for discretionary spending.

In a competitive credit market, 80% are willing to switch lenders for better rates. As LLMs simplify loan comparisons and agents monitor and alert customers to better deals, switching could become more automated and frequent – resulting in increased refinancing and balance transfers.

Base: 480 credit-active, digitally literate Indian consumers
Source: Experian research conducted by Forrester Consulting, July 2026

This combination of confidence and financial pressure creates an important context for the Agentic Consumer. Consumers do not necessarily need more information; they increasingly need help making sense of information and translating it into appropriate action. An intelligent financial assistant could potentially help a consumer understand commitments, compare alternatives, anticipate repayment pressure and identify relevant actions before a financial issue becomes more difficult to manage.




Many consumers are willing to consider additional credit, but their decisions are increasingly shaped by interest rates. The research also shows that consumers are highly open to switching providers when better rates or offers are available, which is a known consumer market practice in India. As AI tools make product comparison easier, faster and more personalised, financial services providers may face a more competitive credit market.

B. AI is Becoming a New Decision-Support Layer

Consumers are increasingly using LLMs to help them understand and compare financial and purchasing decisions. These tools are becoming a new interpretation layer between consumers and providers, helping people make sense of financial terminology, product options, loan comparisons and credit information. Consumers are already using them not only to make informed financial decisions, but they also show a clear preference for tools that are connected to trusted providers and designed specifically for financial services.

Top LLM Use Cases

"Yes, I have used a general-purpose LLM for the following during the past year."



Consumers are more trusting of LLMs provided by a global tech company (91% very/somewhat trusting) or a global payment network (90%) than they are of a FinTech lender / digital bank (78%) or their preferred social media provider (77%)

Base: 480 credit-active, digitally literate Indian consumers

Source: Experian research conducted by Forrester Consulting, July 2026



The Forrester survey shows consumers using general-purpose LLMs for a range of practical financial activities:

- 45% have used them to understand financial terms such as interest rate, fees or credit scores
- 43% choose to use LLMs to compare loan options from different providers
- 49% use them to help manage a personal budget or track spending

This represents an important evolution in the role of digital technology. Traditional digital channels made financial information easier to access. AI increasingly makes that information easier to interpret and compare. Consumers can potentially ask questions in their own language, simplify complex terminology, examine trade-offs between products and bring information from different providers into a common frame of reference. AI therefore begins to sit between the consumer and the provider as an additional interpretation layer.

The research also suggests that the source of this intelligence matters. The opportunity is therefore not simply to provide consumers with more AI. It is to provide trusted AI that combines intelligence with reliable information and appropriate financial context.

C. From AI Assistance to Agentic Action

The next stage is the movement from an AI that helps consumers understand a decision towards an agent that can potentially help them execute it.



Order based on highest proportion of "strongly agree"

Base: 480 credit-active, digitally literate Indian consumers
Source: Experian research conducted by Forrester Consulting, July 2026



The consumer proposition is compelling.

- 94% agree that an agent could help them compare far more options than they could manually
- 92% believe an agent could help them save money by finding the best prices or rates
- 91% believe an agent could save them time by handling research and purchasing
- 90% agree that an agent could help them avoid missing important details such as fees or contractual terms

These benefits suggest that agentic technology could reduce several forms of friction simultaneously. The value proposition therefore moves beyond automation. It becomes one of greater choices, with lower efforts and more informed decision-making.


What are the most important agent capabilities that would influence consumers' decision to use one?

- 89% say detecting unusual charges across payment methods that may indicate fraud and notifying them immediately
- 89% say proactive alerts when there is a risk of missing a payment or going into overdraft, and suggesting ways to avoid fees
- 89% say keeping track of monthly expenses and monitoring new offers to help switch to better deals

D. But Delegation Has Boundaries

Consumer enthusiasm does not mean unlimited willingness to delegate. The research indicates that consumers become progressively more cautious as the agent moves from research and recommendation into identity, data sharing and financial commitment.

Top potential risks of an AI agent making a purchase or application on your behalf



Order based on highest proportion of "strongly agree"

Base: 480 credit-active, digitally literate Indian consumers
Source: Experian research conducted by Forrester Consulting, July 2026



The survey also highlights some of the risks consumers associate with agentic activity.

- 72% are concerned that an agent could make decisions based on incorrect or outdated information
- 72% that it may favour certain brands or providers
- 71% that personal data could be accessed or misused
- 70% that an agent could be manipulated or impersonated by cybercriminals

The implication is clear: the more an agent can do, the stronger the trust infrastructure around it needs to become. A consumer may be comfortable asking an agent to compare five loan products with relatively limited controls. Allowing the same agent to access financial data, prove identity, submit an application on their behalf or accept an offer introduces materially different requirements around consent, security, transparency and accountability.

Taken together, the findings provide a clear snapshot of the **Agentic Consumer in India** – balancing strong trust and perceived value with financial pressures and a continued desire for control.

India

India shows one of the strongest combinations of trust and value perception in the study. It has the highest trust score for LLM accuracy, and close to a third of respondents would grant their agent semi-autonomy. Indian consumers are also among the most comfortable sharing financial data with an agent during an automated loan application. This suggests strong potential for agent-assisted financial journeys, particularly where AI helps consumers compare options and make faster decisions.

LLM Trust Score



8.1/10

Trust the accuracy of the information provided by LLMs

Average: **6.8**

Max agent spend


12%

Feel comfortable delegating \$250 of agentic purchases for online goods

Agent readiness

Optimisers

Efficiency seekers, results focused



Sceptics

Cautious and concerned, need reassurance

Pragmatists

Balanced and open, value clear benefits

Agent delegation readiness

Applying for a loan or a credit card

36%

No autonomy

Recommendations only (no actions taken)

32%

Semi-autonomy

Can act only after approval

23%

Conditional autonomy

Can act if certain rules are met

9%

Full autonomy

Can always act independently

Agent comfort

74%

Would feel comfortable sharing financial data with an agent for automated loan applications

#1 Expected benefit of agent use



94%

"Help compare far more options than they could manually"

#1 Concern of agent use



72%


"Overlooking important details"

Base: 480 credit-active, digitally literate Indian consumers

Source: Experian research conducted by Forrester Consulting, July 2026



Credit health



Base: 480 credit-active, digitally literate Indian consumers

Source: Experian research conducted by Forrester Consulting, July 2026

Chapter 2



The Role of Agentic AI in the Indian Credit Landscape

From the Agentic Consumer to the Intelligent Lender

The Agentic Consumer therefore represents one side of a potentially significant transformation in credit.

Consumers may increasingly arrive at a financial decision supported by an agent that understands their objective, has compared available propositions and, with permission, can organise the information needed to progress an application. But a consumer-side agent cannot independently determine how a regulated lender should assess affordability, verify identity, detect fraud, price risk or make a credit decision. Those responsibilities remain with the financial institution.

The next evolution of the credit journey therefore depends on **connecting consumer intelligence with lender intelligence**.

On one side is the consumer and their AI agents — bringing intent, preferences, permissions and relevant information. On the other is the lender — bringing identity and fraud controls, credit data, analytics, policy, decisioning and regulatory accountability. Between the two sits the trust layer required to allow information and actions to move safely from one environment to the other.

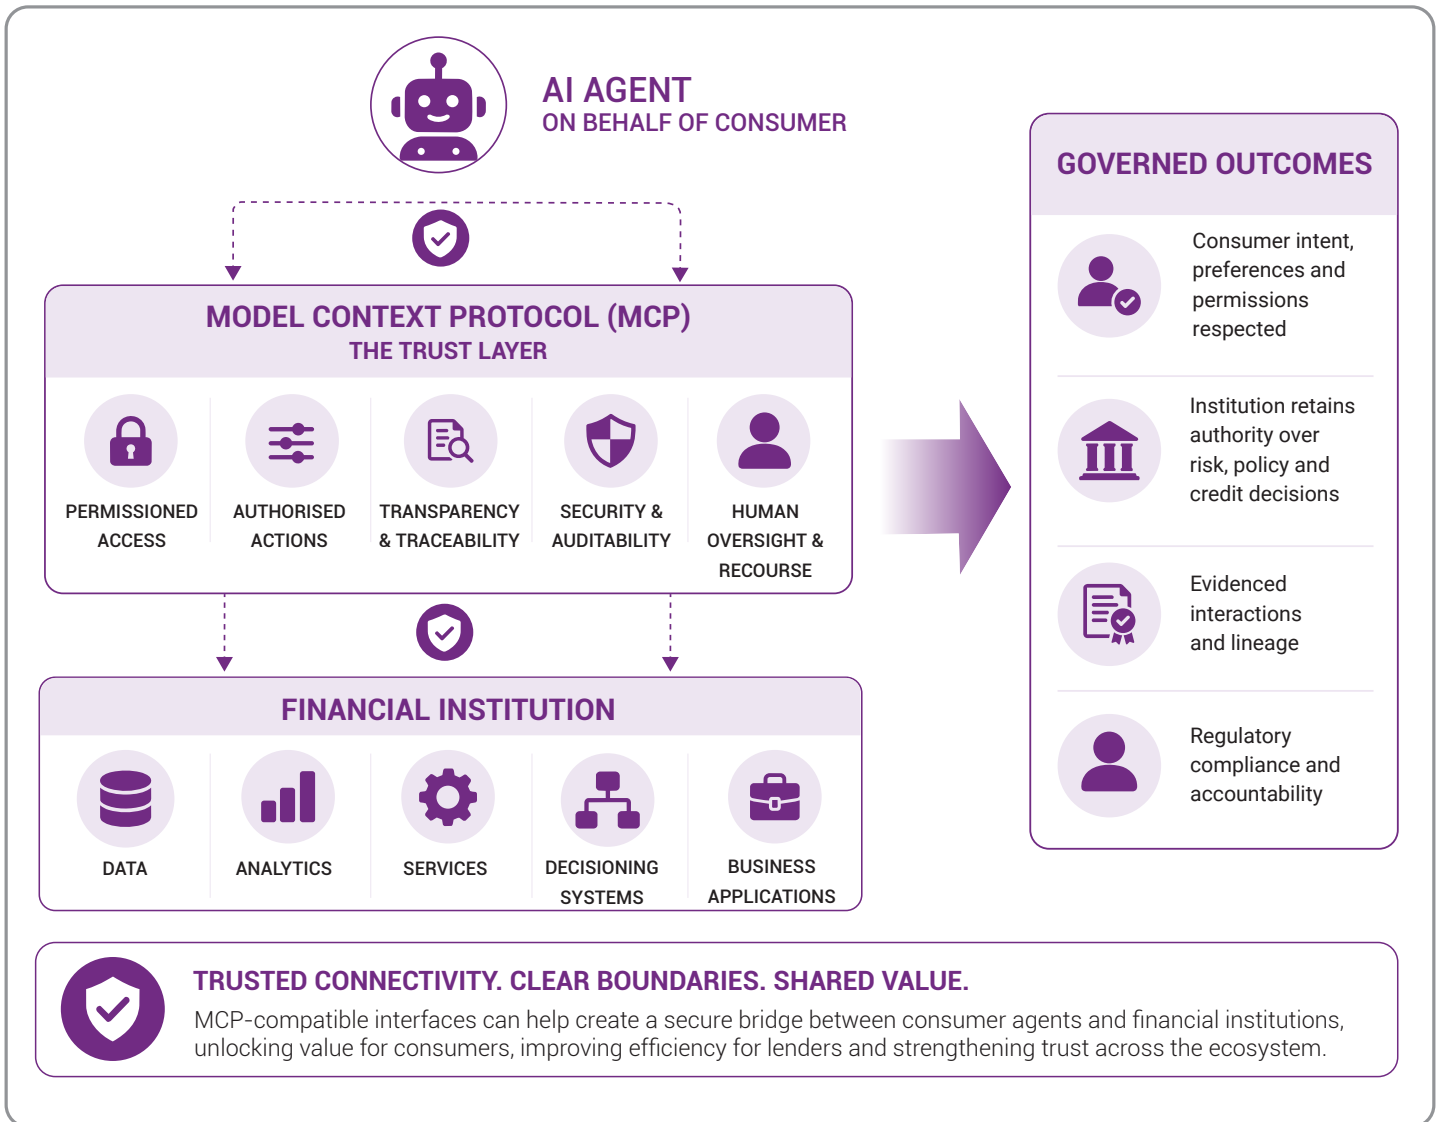
Model Context Protocol: Building the Trust Layer for Agentic Finance

As AI agents become more capable, the future of financial services will depend not only on intelligence, but also on how that intelligence securely connects to data, systems, and institutions. Consumers may increasingly rely on agents to research products, compare lenders, prepare applications, and manage financial decisions. To make these interactions reliable at scale, the industry will require a standardised framework that governs how AI systems access information and perform actions.

This is where **Model Context Protocol (MCP)** becomes important. MCP provides a common framework that enables AI agents to securely connect with enterprise data, analytical services, business applications, and decisioning systems. Rather than creating isolated AI experiences, it allows intelligence to move across the ecosystem while maintaining permissions, transparency, security, and auditability.

In financial services, MCP has the potential to act as the trust layer between consumers, AI agents, and lenders, ensuring that information is exchanged safely, actions are properly authorised, and decisions remain governed. As agentic ecosystems evolve, standards such as MCP could become a critical foundation for enabling innovation without compromising trust.





Evolution in India's Credit Landscape

As financial institutions look to deliver faster, more personalised and efficient customer experiences, AI is beginning to evolve from a largely analytical and assistive capability into one that can potentially orchestrate tasks, make contextual recommendations and enable actions across complex workflows.


Agentic AI represents an important next step in this evolution. Unlike conventional AI applications that typically respond to pre-defined inputs or perform individual tasks, agentic systems can be designed to pursue defined objectives, interact with multiple data and technology environments, coordinate sequences of activities and adapt their actions based on changing information - all within established governance and human-oversight frameworks.

India presents a particularly relevant environment for this transition. **The scale of its consumer market, maturity of its digital ecosystem, growth of digital lending and fintech participation, and continued expansion of formal credit create significant opportunities for AI-enabled innovation.** At the same time, the complexity of credit decisions and the importance of consumer outcomes mean that adoption must be accompanied by robust approaches to data governance, transparency, consent, explainability and human oversight.

How Agentic AI Capabilities Are Taking Shape in Indian Lending

The emergence of Agentic AI in India can be viewed as the next stage in the evolution of digital lending. Over the past several years, lenders have invested significantly in digital onboarding, automated decisioning, machine-learning-based risk assessment, fraud analytics and conversational AI. Agentic AI has the potential to build upon these foundations by connecting individual capabilities into more autonomous, goal-oriented workflows that can interpret information, determine the next best action, interact with systems and escalate cases to humans where required.

The evolution towards Agentic AI in Indian lending can be understood across three stages:



This progression is particularly relevant to India because of the scale and digital nature of its financial ecosystem. The opportunity is therefore not simply to automate existing processes, but to re-design credit journeys around intelligent systems capable of converting data and analytics into co-ordinated action. Within Indian lending, this evolution is becoming particularly relevant in delivering impact across the entire credit cycle.

From Digital Rails to Intelligent Action

The Emerging Agentic AI Stack for Indian Credit

Digital public infrastructure

Trusted digital rails for identity, data, documents, payments and consent



AADHAAR

Digital Identity



UPI

Real-time
Payments



DIGILOCKER

Digital
Documents



ACCOUNT
AGGREGATOR

Consent-based
Data Sharing



ULI

Interoperable
Lending



OUTCOME

Trusted Infrastructure
Enabling Connected
Credit Journeys

Intelligence infrastructure

Turning data into insights, predictions and understanding



ML

Predict risk &
behaviour




ADVANCED
ANALYTICS

Discover patterns &
opportunities



GEN AI

Interpret, summarise
& generate



LANGUAGE AI

Understand,
converse & engage
(multi-language)



OUTCOME

Intelligence That
Makes Decisions
Smarter and More
Contextual

Agentic AI – The Action Layer

Orchestrating actions to achieve objectives within governance and guardrails




REASON

Understand
context &
objectives



PLAN

Determine best
course of
action




ORCHESTRATE

Coordinate
data, models,
rules & tools



ACT

Execute or
recommend
actions



LEARN

Learn from
outcomes &
improve




OUTCOME

Autonomous Action.
Continuous Learning.
Responsible
Outcomes.




Three Layers. One Purpose: Transforming data and infrastructure into intelligent action – responsibly, at scale, for better credit outcomes.

Within Indian lending, this evolution is becoming particularly relevant in delivering impact across the entire credit cycle, as demonstrated in the below diagram.

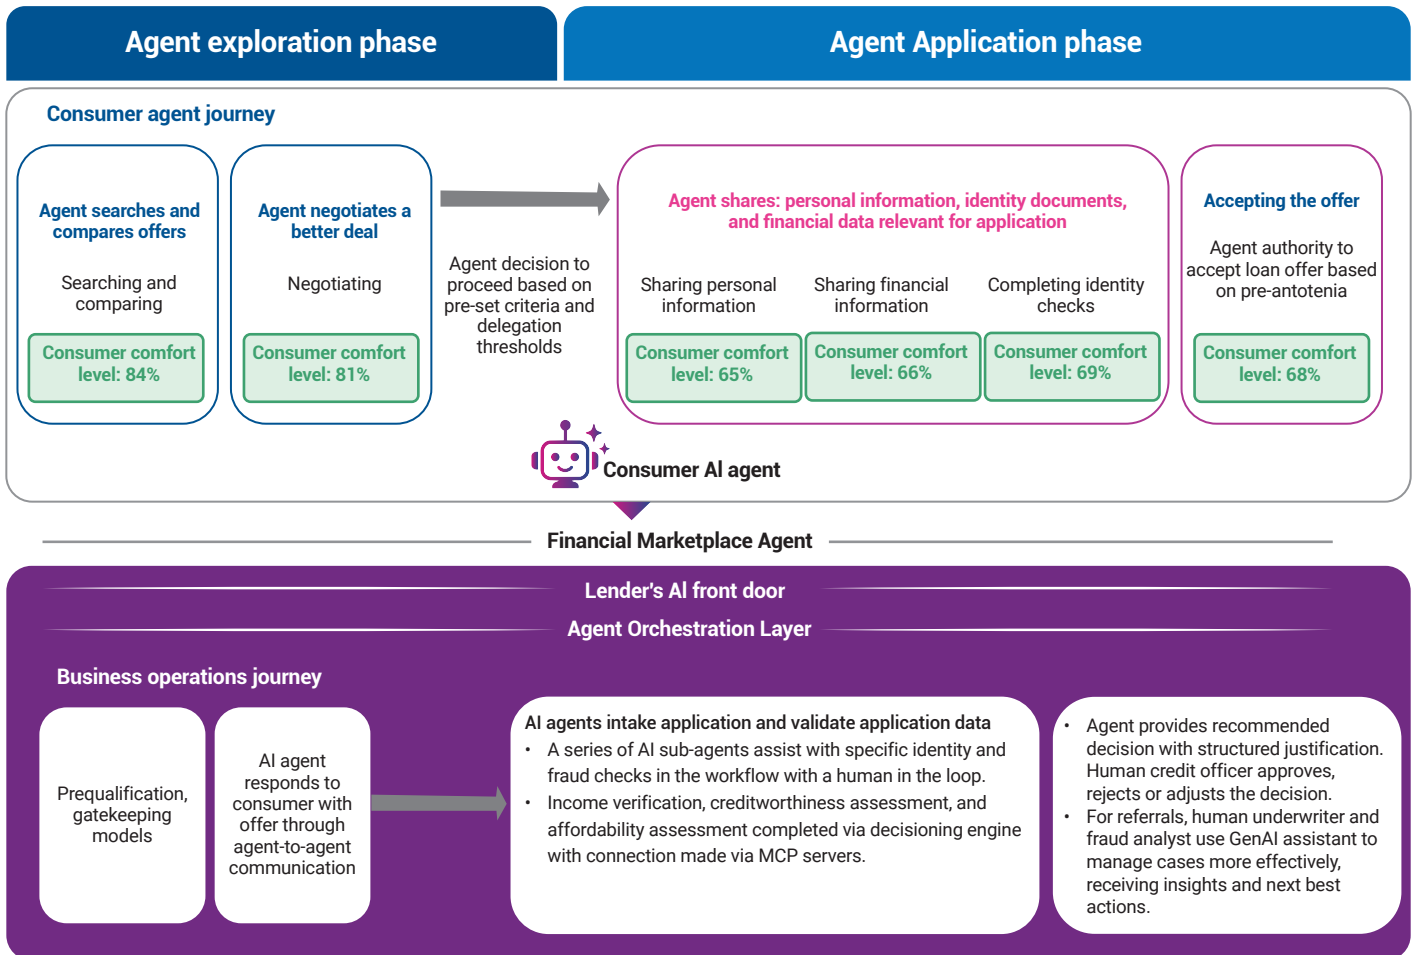


Role of One Platform

Carry connected intelligence and evidence across the journey without becoming the credit authority



Personal Lending brings this connected-platform model into sharp focus. Its high digital maturity, large volumes, strong fintech participation and rapidly changing risk profiles make it an ideal proving ground for Agentic AI. The diagram represents how Agentic AI can transform the personal lending journey by enabling autonomous, coordinated interactions between borrowers and lenders and financial data platforms – from offer discovery and negotiation through application, verification and offer acceptance.



Chapter 3




The Personal Loan Story – Focus
on Fintech-led growth

Digital credit has scaled. The next challenge is ensuring that risk intelligence can keep pace.

Agentic AI can support acquisition, verification, underwriting, monitoring and collections across retail credit. Personal lending is the most practical first test: journeys are already digital, fintech participation is high, transaction volumes are large, unsecured risk can change quickly and outcomes become visible relatively fast.

Scale makes the case. As per Experian data, in FY'26, 51.69 crore loans worth ₹97.7 lakh crore were sourced by lenders across the Indian retail credit market. Personal cash lending contributed 16.23 crore loans worth ₹11.4 lakh crore - 31% of volume but 12% of value. In FY'27 Q1, it represented 35% of volume and 12% of value. Digital-native lenders have already shown that acquisition, evidence assembly and decisioning can operate at scale; the next step is governing orchestration and continuous risk intelligence.

Why personal lending is the first test case for Agentic AI





The proving-ground logic

Personal lending does not dominate credit value, but its digital maturity, operating scale and faster feedback loops make both opportunity and control gaps visible quickly.

Volume sits in small tickets; value sits in large tickets

Ticket size creates two different operating models. In FY26, loans below ₹1 lakh were 87.7% of new accounts but only 16.0% of sanction value; loans above ₹5 lakh were 3.8% of count but 59.1% of value. Small-ticket economics therefore depend on low unit cost, straight-through evidence and fraud controls, while larger exposures require deeper affordability, exposure and pricing assessment.







Ticket segment	FY26 count share	FY26 value share	Primary decision need
Below ₹1 lakh	87.7%	16.0%	Straight-through evidence orchestration, fraud controls, low unit cost and scalable servicing.
₹1 lakh-₹5 lakh	8.5%	24.9%	Balance automation with income, affordability and policy validation; refer genuine exceptions.
Above ₹5 lakh	3.8%	59.1%	Deeper evidence, exposure and pricing discipline, plus ongoing portfolio monitoring.

Interpretation

Ticket size determines journey economics and the depth of evidence required; it should not be treated as a proxy for credit quality.

Fintech-led origination has shifted the operating model

In FY27 YTD, Fintech NBFCs held about 75% of the sanction count for loans below ₹1 lakh and 51% of sanction amount; NBFCs held about 23% of count and 41% of amount. Together they accounted for 97.6% of count and 92.7% of amount, making small-ticket personal lending predominantly a non-bank operating model.





ROI = Rest of Industry





Scale has improved, but risk performance still differs by lender and cohort



Fintech NBFC Net 30+ fell from 11.3% in Jun'23 to 6.5% in Jun'26 and Net 90+ from 5.7% to 2.7%, yet both remained the highest among the categories shown. Results must be read with portfolio mix, vintage, acquisition channel and growth. Risk-adjusted scale therefore requires cohort-level early warning rather than a point-in-time approval view



Growth is broadening; workflows must become context-aware


Both personal loan segments show a strong rebound in FY27 Q1, with growth accelerating sharply across Tier-1, 2, 3 cities, while Tier-4 growth is comparatively weaker. Despite this, Tier-4 continues to dominate market share (~54% for <₹1L and 50% for ₹40K–80K), although its share has declined from FY26 Q1 as Tier 1–3 cities gained share.





What this means for Agentic AI

One workflow cannot serve every ticket, channel and cohort. Personal lending needs context-aware evidence assembly, risk-proportionate verification, continuous portfolio intelligence, governed human referral and closed-loop learning.



The bridge to governed decision velocity

The personal-lending story establishes three facts.

- Operational volume and portfolio value sit in different ticket-size segments.
- Fintech NBFCs and NBFCs have become the primary operating engines of small-ticket credit.
- Portfolio performance and geographic growth can diverge materially across lender categories and cohorts.

The challenge is therefore no longer only to detect risk. It is to convert changing signals into validated models, approved strategies, controlled production decisions and monitored outcomes before the market moves again. Risk insight creates value only when it can be translated into governed action at decision speed.



Chapter 4



Trust at Scale – Data Governance and
Regulatory Checks for Agentic AI

The six foundations of consumer trust in AI agents

1. Identity, verification, liability

- ✓ Secure authentication before sensitive actions, using biometric or ID verification or fingerprint
- ✓ Clear legal liability if AI agent makes an error that costs me money
- ✓ Secure and verifiable link between you and the agent using identity and fraud-prevention checks



Showing all responses where respondents ranked 'very important' and 'important' above 80%



2. Approval & Reversibility

- ✓ Personally approve purchases/applications before the agents complete them
- ✓ The ability to override or cancel any decision my agent makes within a certain time frame



3. Financial controls

- ✓ Setting up specific spending limits
- ✓ Agent purchases/applications protected by the same dispute and liability rules that apply to traditional payments
- ✓ Explicit approval of any decision that creates a financial commitment



4. Transparency & Explainability

- ✓ A step-by-step breakdown of all the reasoning that the agent did to reach its decision
- ✓ Seeing information sources and data that the agent used to make its recommendation



5. Trusted payment infrastructure

- ✓ Ability for the agent to use trusted, secure payment rails, such as a credit card or payment provider
- ✓ Connecting my agent to a secure digital wallet so it can make payments without directly accessing my bank or card details



6. Data & Privacy

- ✓ Ability to control the data the agent stores, how long it stores it, and how it uses it




Consumer trust in AI depends on clear safeguards around identity, transparency, control, privacy and accountability. These considerations are increasingly shaping how regulators and institutions approach AI governance across global markets. AI can make credit journeys faster and more responsive, but the benefits only hold if the institution can explain the data used, the decision reached and the accountability behind the outcome. This is why the governance conversation is moving from broad principles to controls that sit inside the lending journey itself.

That shift is visible across global markets. Regulators may use different models, but most are converging around the same practical questions: what data is being used, how the decision can be traced, what authority an AI system has, and when human oversight is required.

Global best practices

Across markets, AI governance is moving from broad principles towards more practical controls around data, decision-making, human oversight and accountability. The approaches differ, but the common direction is towards stronger controls as the impact of an AI-supported decision increases.



Across these approaches, the common principle is straightforward: institutions need to understand the data being used, keep decisions traceable, apply controls according to risk and retain clear accountability. With Agentic AI, this extends to defining what the system is allowed to access and do.

For lenders operating across markets, the implication is practical: a common control framework should govern data, models, agents, human authority and evidence. Jurisdiction-specific requirements can then be layered onto that base depending on the use case and level of risk.

This global direction connects directly with the Indian banking context. The case for using AI is strong, but it is not a case for unconstrained automation. It is a case for responsible enablement: AI can support credit delivery, service, inclusion, efficiency and fraud prevention, while the institution keeps control over data, decisions and accountability.

The Case for AI in Indian Banking

RBI's position is clear: AI should be responsibly harnessed, not simply treated as a risk to be contained. RBI Governor, Sanjay Malhotra, highlighted has five areas where AI can strengthen Indian banking:

1. **Credit and risk** - AI can use alternative data such as cash flows, GST filings and utility payments to improve assessment of borrowers with limited traditional credit history. It can also help banks identify emerging credit and liquidity risks earlier.
2. **Customer service** - Used to augment rather than replace human judgement, AI can help relationship managers serve customers more effectively and support faster grievance redressal and personalised financial guidance.
3. **Financial inclusion** - Indian-language voice interfaces can make banking more accessible, while predictive models can identify borrowers showing early signs of financial stress. AI can therefore support both wider access to credit and earlier intervention.
4. **Operational efficiency** - AI can automate document processing, reconciliation, regulatory reporting and other routine activities, reducing manual errors and allowing employees to focus on work requiring judgement.
5. **Fraud prevention** - As the Governor noted, "Fraud today moves at the speed of an API call." Machine-learning models can continuously analyse transaction patterns and identify anomalies in real time.

The approach, however, cannot be uniform. Each bank needs to shape its AI strategy around its customers, risk appetite and ability to govern the technology it deploys.


Source: Sanjay Malhotra, Governor, Reserve Bank of India, "Winning in the AI Era: The New Playbook for Indian Banks", FIBAC 2026, Mumbai, 11 August 2026.

Responsible AI: moving from principles to practice

The approaches discussed above show that responsible AI is increasingly being defined not only by regulation, but by how organisations put principles such as transparency, fairness and accountability into practice. **Experian's Responsible AI framework builds on this by setting out seven principles that provide a practical foundation for developing and using AI responsibly.** Together, they cover the quality and performance of AI systems, the data behind them, how risks and bias are managed, and whether decisions remain secure, explainable and accountable. These principles become particularly relevant as AI, and increasingly Agentic AI, moves closer to customer interactions and credit decisioning.

Experian's Responsible AI Framework

Seven principles of Responsible AI



For lenders, Responsible AI means ensuring reliable models, suitable data, controlled risks, security, explainability, privacy and fairness. With Agentic AI, these controls must also extend to the data accessed, tools used and actions taken. As adoption grows, strong data foundations and human oversight will remain central to maintaining trust. Ultimately, Responsible AI is a process, not a destination.


Chapter 5



Harnessing the Power of AI
through Platforms

Risk insights create value only when they can be translated into governed action at the speed at which portfolios change

The preceding personal-lending story showed how risk is evolving as fintech-led origination scales: customer leverage can change quickly, portfolio performance can diverge across cohorts and channels, and early-warning signals can appear before headline delinquency measures move. The next question is therefore operational, not analytical: how quickly can a lender convert a new insight into a defensible model, an approved strategy, a production decision and a monitored outcome?




The central proposition is simple: automate the mechanics of the model lifecycle, not the accountability of model judgement.

From evolving risk to decision velocity

The challenge is no longer only detecting risk. It is operationalising insight before the market moves again.

Personal-lending portfolios now generate a continuous stream of signals: bureau changes, utilisation trends, new-credit searches, repayment behaviour, vintage performance, channel mix and affordability indicators. These signals can reveal both opportunity and stress, but they do not create value by themselves. Value emerges when the institution can translate them into consistent decisions and interventions.

In many organisations, that translation remains slow. A business question is converted into a modelling brief; data is located and prepared; a model is built; documentation is assembled; validation and model-risk review take place; deployment is handed to another team; inference is operationalised; and monitoring is established later. Each step may be individually strong, yet the handoffs between them absorb time and break context.



The decision-automation gap

Digital lenders may be able to identify emerging risks quickly yet still require weeks to move a model or strategy from analysis to controlled production. Agentic orchestration is valuable when it reduces that gap without weakening validation, explainability or oversight.

Why a connected lifecycle matters

- The same business intent and data context can travel across development, validation, deployment and monitoring instead of being recreated at each boundary
- Specialist agents can prepare evidence, documentation and operational artefacts while experts retain ownership of assumptions, model choice and policy decisions
- Monitoring outcomes can return directly to the analytical environment, shortening the path from observed deterioration to controlled refinement
- A common platform can standardise the evidence trail across products, channels and model teams, making scale more auditable rather than less controlled

A useful design principle

The objective is not a self-directing “model factory”. It is a governed operating environment in which users express an outcome, agents coordinate approved tasks, and accountable teams remain responsible for every material decision.



From disconnected tools to one intelligent experience

The platform shift is from manual movement between systems to a shared, agent-orchestrated context.


Lifecycle moment	Disconnected by default	Agent-powered evolution
Analysis	Notebooks and code operate separately from downstream deployment and governance context.	A user states the intended outcome; the assistant invokes the relevant analytical agents and approved tools.
Feature engineering	Data preparation and feature work require separate movement, reconciliation and version management.	Feature preparation is coordinated within the same governed context, with lineage preserved.
Validation	Model cards, statistical tests, fairness evidence and policy mapping are assembled manually.	Specialist agents draft validation evidence and route it to model-risk and compliance reviewers.
Deployment	Artefacts are transferred to engineering or operations through manual handoffs.	Packaging, registration, promotion and inference are orchestrated through Ascend Ops.
Monitoring	Health and drift are reviewed through separate dashboards and periodic processes.	Performance monitoring remains connected to the model version and can trigger a governed review or retraining workflow.

The differentiated value is not the conversational interface alone. It is the governed intelligence beneath it: lender data, Experian credit information, trended attributes, identity and matching capabilities, affordability signals, fraud indicators, model assets and decision strategies. The agent layer becomes useful because it can invoke these assets in a controlled sequence and preserve the evidence behind each step.

A unified platform across Sandbox, Ops and Decisioning

A common orchestration layer connects build, deployment, model-risk governance and feedback. The platform vision brings the lifecycle into one connected environment. Specialist agents can be invoked conversationally within a product or headlessly through approved APIs, batch workflows or MCP-compatible orchestration. The user expresses the required outcome; the platform orchestrator sequences the appropriate agents and preserves the context across each environment.

Ascend Platform



What becomes different

- Data can remain within the governed environment while agents operate on permitted context, reducing uncontrolled file movement and duplicated extracts
- Validation evidence is generated alongside the model rather than reconstructed after development is complete
- The deployed model remains a living system: inference, health, drift, troubleshooting and model-risk evidence stay connected to the registered version

Why this matters for fintech risk

When portfolio risk changes quickly, the institution needs to test and operationalise a response without rebuilding the process around each new insight. A connected lifecycle makes speed repeatable and governable.

Illustrative use case: the Responsible Growth Agent

Growth and customer protection can be treated as one decision problem, not two competing agendas.

Consider a personal-lending portfolio in which overall growth remains attractive, but risk is evolving unevenly. Some customers show stable repayment behaviour and increasing capacity. Others appear positive but lack a current affordability view. A third group is beginning to show rising leverage, missed payments or repeated credit searches. A blanket growth campaign would ignore these differences; a blanket tightening would leave good opportunity unused.

The business question

Who can we responsibly lend more to, who needs an affordability check first, and who needs support rather than more credit?

A risk or portfolio manager states the required outcome in business language. Experian Assistant then coordinates the relevant agents, data and decision assets. The output is not an unconstrained AI recommendation; it is a segmented, explainable proposal that can be validated, approved and deployed through governed strategies.

Segment	Management intent	What the agent sees	Governed action
GREEN	Responsible growth	Stable or improving behaviour; low external risk; manageable utilisation; affordability already established.	Consider a policy-approved top-up or limit increase, subject to eligibility, pricing and exposure rules.
AMBER	Verify first	Positive signals, but current affordability is not sufficiently confirmed.	Invite purpose-specific consent for an affordability assessment, for example through approved Account Aggregator or bank-data analysis, before any offer.
RED	Protect and support	Rising utilisation, recent missed payments or heavy new-credit search indicating early-warning risk.	Do not extend new credit; route to an appropriate support, financial-health or servicing pathway.



The AI angle

Agentic AI changes the experience by allowing the user to start with the outcome rather than the mechanics. The assistant can identify the data required, invoke specialist agents, organise the evidence and explain the proposed segmentation in plain language. But the growth offer, affordability requirement and support decision remain anchored to approved models, policy rules and human sign-off.


What the agent cannot do

It cannot redefine delinquency, invent an affordability threshold, suppress an adverse signal, create an unapproved offer, waive model-risk review or deploy a strategy without the required approval.

From the model lifecycle to the loan lifecycle

The personal-lending analysis established why risk intelligence must evolve: borrower leverage, repayment capacity, channel performance and portfolio quality can change faster than traditional policy and model-change cycles. This chapter has addressed the institutional response - how a lender can convert those changing signals into validated models, approved strategies, controlled production decisions and monitored outcomes.

A governed model lifecycle is not an end in itself. Its value is realised when approved data, models and strategies are invoked at the right moment in a customer or portfolio journey. The same intelligence that helps a lender identify risk can support pre-approval, evidence collection, affordability assessment, pricing, referral, fulfilment, early warning and customer treatment.



Chapter 6

The Road Ahead

From the Agentic Consumer to the Agent-Ready Credit Enterprise

Consumers are beginning to use AI not only to interpret financial products, but to define goals, compare options and move from intent to action. In an agent-mediated market, the first question may be 'What can I responsibly afford?' rather than 'Which loan should I apply for?' The customer's agent may organise context and permission, discover providers and compare structured propositions; the regulated institution still owns affordability, fraud prevention, policy, pricing, disclosure and the final credit decision.

This changes the basis of competition. Brand visibility and marketplace placement may no longer secure the first interaction. Advantage shifts toward institutions that are discoverable and callable by authorised agents, can return transparent propositions at machine speed, and can preserve evidence and accountability from request to outcome.

What Changes in an Agent-Enabled Credit Market

Agentic AI shifts the lending journey from provider-led experience to an intelligent agent-led ecosystem. MCP-compatible interfaces and APIs can connect agents to approved capabilities while trusted intelligence and evidence layers enable lenders to respond with speed, relevance and control. Decisioning platforms, like Experian Ascend, can translate these insights into policy-aligned actions, with human oversight retained wherever judgement and accountability matter.

The shift has four important implications.

- **Access shifts to AI agents:** The consumer's objective and context may sit with a personal agent rather than a provider-owned screen.
- **Distribution becomes machine-to-machine:** MCP-compatible tools and APIs can expose bounded discovery, consent, evidence, decision-request, offer and status services.
- **Decision speed becomes strategic:** Lenders must translate permitted data into responsible, explainable offers in seconds without weakening policy or human referral.
- **Ecosystem roles are redefined:** Banks and NBFCs become agent-ready providers; fintechs can build agents, orchestrate networks or supply specialist tools; Experian anchors trusted intelligence.

The marketplace function may not disappear; it will evolve.

Aggregation, normalisation, trust, routing and fulfilment remain valuable even if the visible comparison interface becomes less central. In an agent-led ecosystem, the winning proposition will therefore be less about owning the customer screen and more about becoming discoverable, trusted, interoperable and ready to fulfil decision-making at machine speed.



Leadership Agenda

Agentic AI should be viewed not simply as a technology initiative, but as a combined customer-access, growth, risk and operating-model transformation. The objective is not maximum autonomy. It is faster, more relevant and more explainable decision-making within clear institutional boundaries.

Five priorities for responsible scale

- **Become agent-ready, not only digital:** Expose accurate products, scoped agent access, low-latency prequalification, structured offers and human escalation.
- **Separate orchestration from credit authority:** Agents may interpret intent and invoke approved tools; governed models, strategy and people remain authoritative.
- **Build one governed intelligence foundation:** Connect permitted data, features, models, validation, decision assets, monitoring and evidence.
- **Embed trust in every call:** Carry purpose, consent, identity, least privilege, lineage, reason codes, security, fallback and customer challenge through the workflow.
- **Scale by evidence and continuous learning:** Measure customer, growth, risk, operations and control outcomes together, then return production evidence to governed review.

An Illustrative Path to Scale

- **0-6 Months:** Map agent journeys and authority. Inventory data and tools. Pilot two or three bounded services with explicit consent, escalation and evidence.
- **6-18 Months:** Standardise MCP/API services, integrate Ascend and decisioning, preserve context across the journey, and strengthen monitoring and rollback.
- **18-36 Months:** Reuse decision assets across products, connect origination to monitoring and collections, and operate enterprise governance dashboards.

The Road Ahead

The next era of lending will not be won by the institution with the largest number of AI pilots or the greatest autonomy. It will be won by the institution that can serve authorised agents quickly, transparently and responsibly - while keeping data, decisions and accountability governed.

Agentic AI creates movement. Enterprise intelligence creates coherence. Governance creates legitimacy.



Disclaimer:

The report ("Unlocking Enterprise Intelligence: A Journey of Innovation and Transformation with Agentic AI") released in September 2026 is created by Experian Credit Information Company of India Private Ltd. ("Experian" on an 'as-is' basis and is based on information provided by lenders/credit institutions and publicly available data.

Experian cannot give any guarantee or warranty that any information in this report is complete, accurate, up-to-date, error-free and/or suitable for any specific purposes. Interpretation of any fields may vary based on information provided by the lenders/credit institutions which may lead to difference in some parameters with actuals as per different lenders/credit institutions. This report is intended to represent the industry to a large extent and is not specific to any lender. This report is intended only for informational and reference purpose and cannot be used for any commercial purpose or to make business decisions. The information contained in this report does not constitute any advice (whether legal or investment) and any user of this report should carry out all necessary analysis and seek necessary guidance that is prudent in its opinion before taking any decision based on the information contained in this report. Experian does not guarantee any particular outcome and is not liable in any way whatsoever, for any inaccuracies, errors, or omissions caused by the information in this report. Experian does not accept any responsibility and/or liability whether direct, indirect, incidental, consequential including error, or discrepancies relating to the report. By accessing or using this report, you acknowledge, accept and understand the terms of this disclaimer with respect to your use or access. No part of this report should be copied, circulated, or published (whether in part or full) without prior written approval of Experian.

This report is solely owned by Experian and all intellectual property rights therein vest with Experian.

© 2026 Experian Credit Information Company of India Private Limited, Inc. All rights reserved.

About Experian

Experian is a global data and technology company, powering opportunities for people and businesses around the world. We help to redefine lending practices, uncover and prevent fraud, simplify healthcare, deliver digital marketing solutions, and gain deeper insights into the automotive market, all using our unique combination of data, analytics and platforms. We also assist millions of people to realise their financial goals and help them to save time and money.

We operate across a range of markets, from financial services to healthcare, automotive, agrifinance, insurance, and many more industry segments.

We invest in talented people and new advanced technologies to unlock the power of data and to innovate. A FTSE 100 Index company listed on the London Stock Exchange (EXPN), we have a team of 25,200 people across 33 countries. Our corporate headquarters are in Dublin, Ireland. Learn more at experianplc.com.





Experian Credit Information Company of India, 5th Floor, East Wing, Tower 3,
Equinox Business Park, LBS Marg, Kurla (West), Mumbai - 400 070, India

T: +91 (0) 22 6818 6760 | **F:** +91 (0) 22 6818 6792 |
india.digitalmarketing@experian.com | www.experian.in/business