MIT Sloan Management Review INDIA

Strategy Shift Forum

Powered by salesforce

August 2025

AI in Action: A Playbook from the Frontlines

Cross-functional frameworks and case studies

Executive Summary

Artificial intelligence has entered a new phase in India. What was once limited to pilots and experiments is now being deployed at scale, reshaping how businesses serve customers, manage operations, and create value. This Playbook, developed from the sessions at the Delhi edition of the Strategy Shift Forum, distills those discussions into five practical plays that leaders can apply immediately.

A recurring theme from the Forum was that AI performs best when its scope is tightly defined. As Neil Thompson of MIT's FutureTech noted, today's systems excel at tasks that would take a human only a few minutes, such as document checks or short queries, but performance drops sharply when problems become more complex, open-ended, and less structured. The message is clear: AI is most powerful when paired with human oversight, not as a replacement for it.

The Forum provided concrete illustrations of this shift. SMS Auto Leasing India Pvt. Ltd reduced customer onboarding from several days to just a few hours by automating KYC checks with guardrails. Apollo 24/7 scaled millions of doctor-patient interactions through its Ask Apollo platform, extending care beyond what individual practitioners could deliver. Indian insurers cut claims processing times from five hours to 30 minutes by adopting hybrid human-AI systems. Each case showed how disciplined deployment, backed by strong oversight, translates directly into measurable results.

The Playbook emphasizes three principles: start with precision by narrowing the scope and expanding only once results are proven; integrate AI into existing workflows rather than building stand-alone tools; and measure success through real business outcomes rather than vanity metrics.

For Indian enterprises, global lessons must be adapted to local realities. Regulatory oversight, multilingual customer bases, UPI-powered digital transactions, and data localization requirements shape both the risks and opportunities of AI adoption.

This Playbook is intended as a roadmap for leaders seeking to transform responsibly, with approaches that are practical, measurable, and scalable.

FIVE PLAYS THAT REINFORCE AI SUCCESS

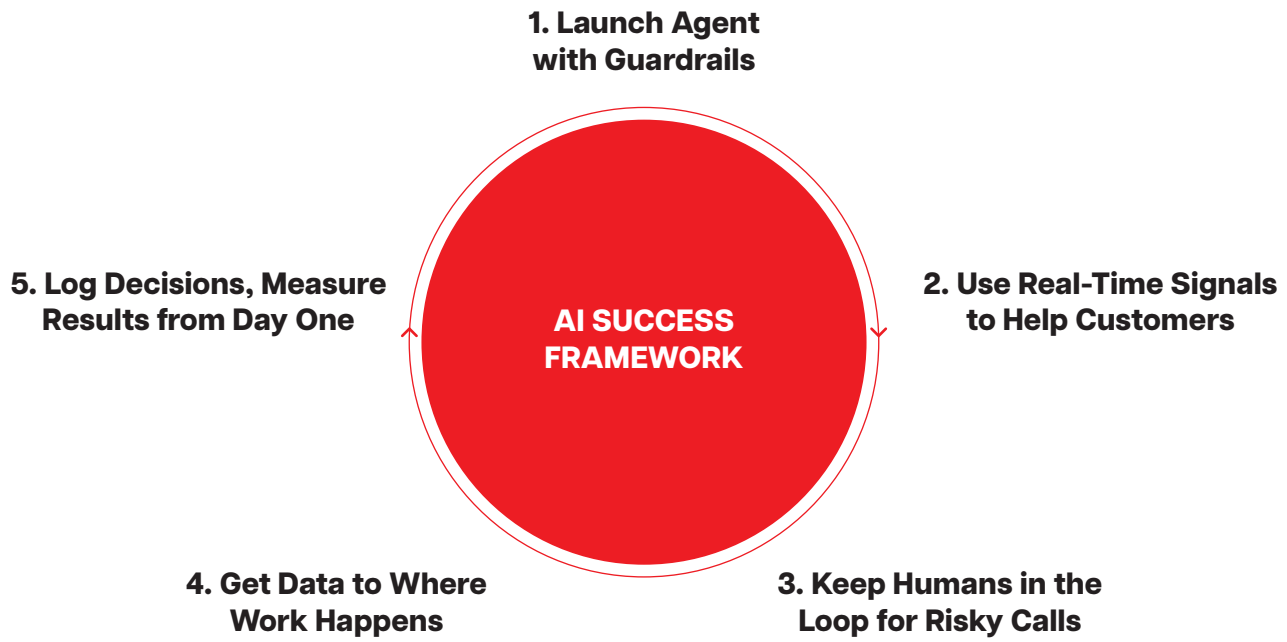

Forum Tracks at a Glance

If you did not attend the Delhi Strategy Shift Forum, start here. The program ran across three tracks that frame the plays in this Playbook. Each explored how AI is reshaping business in practice, not in theory.

Commercial Growth

This track examined how AI is rewriting the rules of selling. Leaders showed how territory planning, propensity-to-buy scoring, and real-time buyer engagement are turning sales from a transactional function into a value-creation engine.

Intelligent Experience

The focus here was on loyalty, personalization, and timely interventions. Discussions highlighted how organizations are using behavioral signals and customer data to trigger relevant actions in the moment, improving satisfaction and long-term engagement.

AI Innovation

This track looked at how enterprises can scale beyond pilots. Speakers emphasized building solid data foundations, embedding innovation into workflows, and maintaining compliance with regulatory and ethical standards while pursuing systemic transformation.

Each of these tracks connects directly to the five plays. Think of the tracks as the themes, and the plays as the practical steps to put those themes into action.

> We are already in the age of autonomous AI. You can flip a switch and let AI control a lot of stuff. The real question is, will it actually do a good job?
>
> **Neil Thompson, MIT FutureTech**

1. Launch Agent With Guardrails

What This Play Does
This play focuses on deploying AI agents in well-defined workflows with built-in escalation and audit protocols. Guardrails ensure compliance, transparency, and trust as agents scale to more complex tasks.

CASE STUDY

At SMAS Auto Leasing India, manual KYC verification once delayed customer approvals by two to three days. Staff had to check identification documents, run verifications across government databases, and manually update records. The firm addressed this bottleneck by deploying an AI-powered leasing assistant.

Rajesh Thanua, General Manager-Technology at SMAS Auto Leasing, said: ***"The one CAM (Credit Appraisal Memorandum) process that used to take two to three days, now my team can do within four to six hours."*** The agent ingests Aadhaar documents, extracts identifiers, and verifies them against national databases in real time. It flags anomalies and escalates complex cases to human reviewers.

Another participant said: ***"By using a low-code environment with plain-English rules, the system is auditable and accessible to non-technical staff."*** This ensured guardrails were in place and compliance standards upheld. The firm has since begun extending the approach into payment cycle analysis and collections, demonstrating scalability.

How to Put It to Work

- Identify a narrow, well-defined workflow for initial deployment
- Define escalation triggers and human review protocols
- Configure agent in a low-code/no-code environment
- Build audit trails and compliance checks
- Monitor performance with real-time dashboards
- Expand to adjacent workflows only after measurable success

How to Measure Progress

Track cycle time reduction, accuracy rates, and escalation frequency. Early pilots should demonstrate 40–60% faster processing while maintaining 95%+ accuracy and fewer than 20% escalations.

Risks to Watch

Over-automation without guardrails risks compliance failures. Lack of transparency in rules and escalation protocols can undermine trust with regulators and customers.

What the Research Says

MIT research highlights that AI performs reliably on bounded tasks but struggles as complexity rises. Success depends on clear workflow boundaries, embedded human oversight, and robust monitoring frameworks.

> There is research showing that people who wear glasses are at least perceived as harder working, more honest, and more intelligent. It's just a silly but scientifically supported example of how small things in our environment can actually have big perceptual and behavioural impacts.
>
> **Patrick Fagan, Capuchin Behavioural Science**

2. Use Real-Time Signals to Help Customers

What This Play Does
This play focuses on shifting from periodic campaigns to real-time interventions that address customer needs in the moment. By listening to live behavioral signals—such as abandoned carts, support queries, or loyalty activity—organizations can trigger context-aware actions that reduce friction and build trust.

CASE STUDY

In the hospitality sector, companies connected loyalty programs with booking engines so that customers could earn or redeem rewards instantly at the point of booking. This eliminated delays and made the loyalty proposition tangible at the very moment of decision. One speaker emphasized: ***"When a customer decides where to stay, if they can see the points right there, it becomes a reason to choose you over someone else."***

A broader consumer ecosystem spanning retail, travel, and finance faced fragmentation. Different business units pushed conflicting offers, creating customer confusion. To solve this, the teams combined offline retail purchases with online app activity, creating a unified customer profile. As one participant noted: ***"The unified profile shapes the customer's next interaction based on real behavior rather than isolated transactions."*** This allowed interventions such as reminding a frequent flyer to use accumulated points at checkout, or suggesting a retail discount to a customer who had recently browsed products online.

For Indian enterprises, the opportunity is even larger. Telecom and retail leaders pointed to UPI transactions and app usage data as rich sources of signals. A speaker observed: ***"Payments data gives you intent in real time, whether someone is topping up a phone plan or buying groceries, and AI can turn that into an offer right then and there."*** These examples illustrate how Indian organizations can leapfrog legacy systems by embedding AI into ubiquitous platforms customers already use daily.

How to Put It to Work

- Map critical journey moments (e.g., booking, cart abandonment, payment)
- Define event triggers and real-time data sources
- Build a unified customer profile across channels
- Establish contextual action rules and escalation pathways
- Pilot interventions with limited segments and measure outcomes
- Scale successful interventions across the ecosystem

How to Measure Progress
Track conversion uplift at the moment of intervention, redemption rates for loyalty offers, reduction in support tickets, and customer satisfaction scores post-intervention.

Risks to Watch
Over-reliance on signals without context can lead to intrusive or irrelevant nudges. Poor data quality or fragmented profiles risk eroding trust.

What the Research Says
Research shows that real-time personalization improves engagement but must be tempered by privacy safeguards. MIT findings emphasize that interventions are most effective when tied directly to observable behavior, not inferred traits.

3. Keep Humans in the Loop for Risky Calls

What This Play Does
This play focuses on striking the right balance between automation and human oversight. AI can handle routine or predictable cases, but when decisions carry financial, legal, or reputational risk, human review remains essential. By designing hybrid systems, organizations ensure accuracy, compliance, and customer trust.

CASE STUDY

An insurer at the Forum shared how they initially attempted full automation of claims processing. The results were unsatisfactory: errors crept in, and customers felt disconnected from the process. The company pivoted to a hybrid approach. ***"The system pre-fills claims based on photos and notes, then an adjuster corrects edge cases on a single screen,"*** explained one participant. This reduced cycle time dramatically from five hours per claim to about 30 minutes.

Another leader underlined the importance of confidence thresholds. ***"You need to define the line where the AI is sure enough to decide on its own, and where it must hand over to a human,"*** they said. By routing only low-confidence cases to adjusters, the company preserved efficiency without sacrificing quality.

The human dimension also mattered. A participant stressed: ***"Our insurance claims examiners were happier because they weren't typing in the same repetitive details. They could focus on tricky cases where their judgment adds value."*** This illustrates how human-in-the-loop systems not only improve outcomes but also enhance employee satisfaction.

How to Put It to Work

- Define confidence thresholds for decision-making
- Categorize processes by complexity and risk
- Create review interfaces with full context for humans
- Implement structured feedback loops to improve AI
- Pilot hybrid workflows before scaling

How to Measure Progress
Track reduction in cycle time per decision, human review rates, error corrections, and customer satisfaction scores.

Risks to Watch
If confidence thresholds are poorly defined, AI may either escalate too much (eroding efficiency) or too little (increasing risk). Employee disengagement can occur if review roles feel redundant.

What the Research Says
MIT research shows that hybrid human-AI collaboration consistently outperforms either full automation or full manual processes in high-stakes contexts. The key is calibrated escalation and transparent decision logging.

4. Get Data to Where Work Happens

What This Play Does
This play emphasizes creating unified data foundations that deliver insights directly into the flow of work. Instead of forcing teams to toggle between systems or wait for offline analysis, data should be seamlessly available in the sales, service, and operations tools employees already use. The goal is to make decisions faster, cleaner, and with less friction.

CASE STUDY

A major enterprise shared how fragmented systems left frontline teams with partial information. Agents often had to pull data from multiple dashboards, slowing service and frustrating customers. One participant recalled: **"*We had support staff wasting five to ten minutes just to piece together a customer's history.*"** This created inefficiencies and missed opportunities to personalize engagement.

To address this, the company built a unified data layer with zero-copy connections across systems. As another leader explained: ***"Teams emphasized a unified data layer with zero-copy connections, allowing agents to work seamlessly within sales, service, marketing, and commerce tools without transferring massive datasets."*** By embedding insights where work happens, employees no longer had to jump across platforms. Decisions became faster, more consistent, and less error-prone.

A third voice stressed governance: ***"It's not just about connecting data, it's about making sure ownership is clear. Who updates, who validates, who is accountable."*** With this structure, the enterprise maintained compliance and trust while improving agility. The initiative demonstrated how embedding data into daily workflows shifts AI from analysis to action.

How to Put It to Work

- Map top 20 critical data events and assign ownership
- Build a unified data layer with zero-copy connections
- Integrate event streaming into frontline workflows
- Set governance rules and access controls
- Deploy real-time dashboards inside existing tools
- Monitor data freshness, quality, and relevance

How to Measure Progress

Key indicators include reduction in time to access complete customer context, fewer errors from manual data entry, and higher first-contact resolution rates in customer support.

Risks to Watch

Without clear governance, a unified data layer can introduce duplication, privacy risks, or confusion about ownership. Another risk is over-engineering—teams may build complex integrations without focusing on actual business outcomes.

What the Research Says

MIT research shows that organizations with integrated data pipelines outperform peers in speed of decision-making. Embedding insights into workflows reduces cognitive load, enabling employees to act on information rather than hunt for it.

> Nobody wants to be put on hold. What if our workforces had no limits, scaled any task 24/7, and delivered quality without compromise?
>
> **Mankiran Chauhan, Salesforce India**

5. Log Decisions, Measure Results from Day One

What This Play Does
This play ensures that every AI deployment is treated like a product from the start, with clear metrics, logging, and continuous monitoring. By capturing inputs, outputs, and outcomes, organizations can improve accuracy, build trust, and tie AI performance directly to business results. The focus is on discipline, measuring what matters and creating a feedback loop.

CASE STUDY

A large enterprise participant described their early challenge: AI agents were deployed, but no one tracked how many conversations closed successfully or what business value was generated. ***"We were flying blind. Agents were running, but we had no logs, no baselines, nothing to prove value,"*** the leader admitted.

They responded by launching an AI command center. ***"One enterprise logged more than 850,000 digital-agent conversations across support and sales, with roughly 30% closed or influenced, and a reported 45% reduction in sales lead time,"*** a speaker said. This visibility helped business and technology leaders agree on next steps and investments.

Another emphasized transparency: ***"Every decision needs a run log that records inputs, prompts, evidence, and outputs. If an auditor asks why an agent made a call, we should be able to replay it."*** By tying each log to a small set of business outcomes, teams ensured AI agents were not optimized for vanity metrics but for results that mattered, including customer satisfaction, revenue growth, and compliance readiness.

This discipline shifted organizational culture. A participant concluded: ***"When we saw week-on-week improvements and knew exactly why, it gave us the confidence to expand use cases."*** The case demonstrated that measurement from day one is the foundation for sustainable scaling.

How to Put It to Work

- Define business outcome metrics for each AI agent
- Create logging architecture to capture inputs and outputs
- Establish baseline performance before deployment
- Set up dashboards for real-time monitoring
- Link every optimization to measurable business results

How to Measure Progress

Track business outcomes such as revenue lift, cost reduction, and customer satisfaction improvements. Monitor cycle time reduction and automation rates, while ensuring compliance and audit readiness.

Risks to Watch

Measuring too many technical metrics can distract from business impact. Lack of transparency in logs creates audit and compliance risks. If feedback loops are not established, AI performance may plateau rather than improve.

What the Research Says

MIT research shows that AI projects with early measurement frameworks are more likely to scale successfully. Organizations that link AI to business KPIs from the start report higher returns and stronger stakeholder confidence.

MEASURING WHAT MATTERS

FOCUS AREA	SUCCESS METRICS
Operational	Cycle time ↓40–60% · Accuracy >95% · Escalations <20%
Business Impact	Revenue ↑20–45% · Cost/transaction ↓30–50% · Customer satisfaction score >4.5/5
Strategic Alignment	1 new use case/month · Adoption >80% · Data errors <5%
Implementation	First agent in 90 days · RoI <6 months · Scale across the organization in a year

Note: Cycle time is the total time taken to complete a process (e.g., resolving a ticket or processing a claim)
RoI: Return on Investment

References

- Renieris, Kiron, Mills, Kleppe. "AI Explainability: How to Avoid Rubber-Stamping Recommendations," 12 Jun 2025.
- Moorman, Hickey. "When AI Investments Pay Off in Marketing," 9 Apr 2024.
- Schrage, Muttreja, Kwan. "How the Wrong KPIs Doom Digital Transformation," 8 Mar 2022.
- Gosline et al. "Nudge Users to Catch Generative AI Errors," 29 May 2024.
- O'Neil, Appel, Tyner-Monroe. "Auditing Algorithmic Risk," 11 Jun 2024.
- Siegel. "What Leaders Should Know About Measuring AI Project Value," 7 Feb 2024.

Resources

Watch highlights from the Delhi edition of the Strategy Shift Forum:

- **Commercial Growth Track**
- **AI Innovation Track**
- **Intelligent Experience Track**

Next stop: **Conrad Bengaluru on 25 September.** Join MIT faculty, global AI experts, and business heads for practical playbooks on scaling AI. For more details, the full agenda, and to request an invitation, **click here.**

MITSloan Management Review INDIA

Strategy Shift Forum

Powered by

salesforce