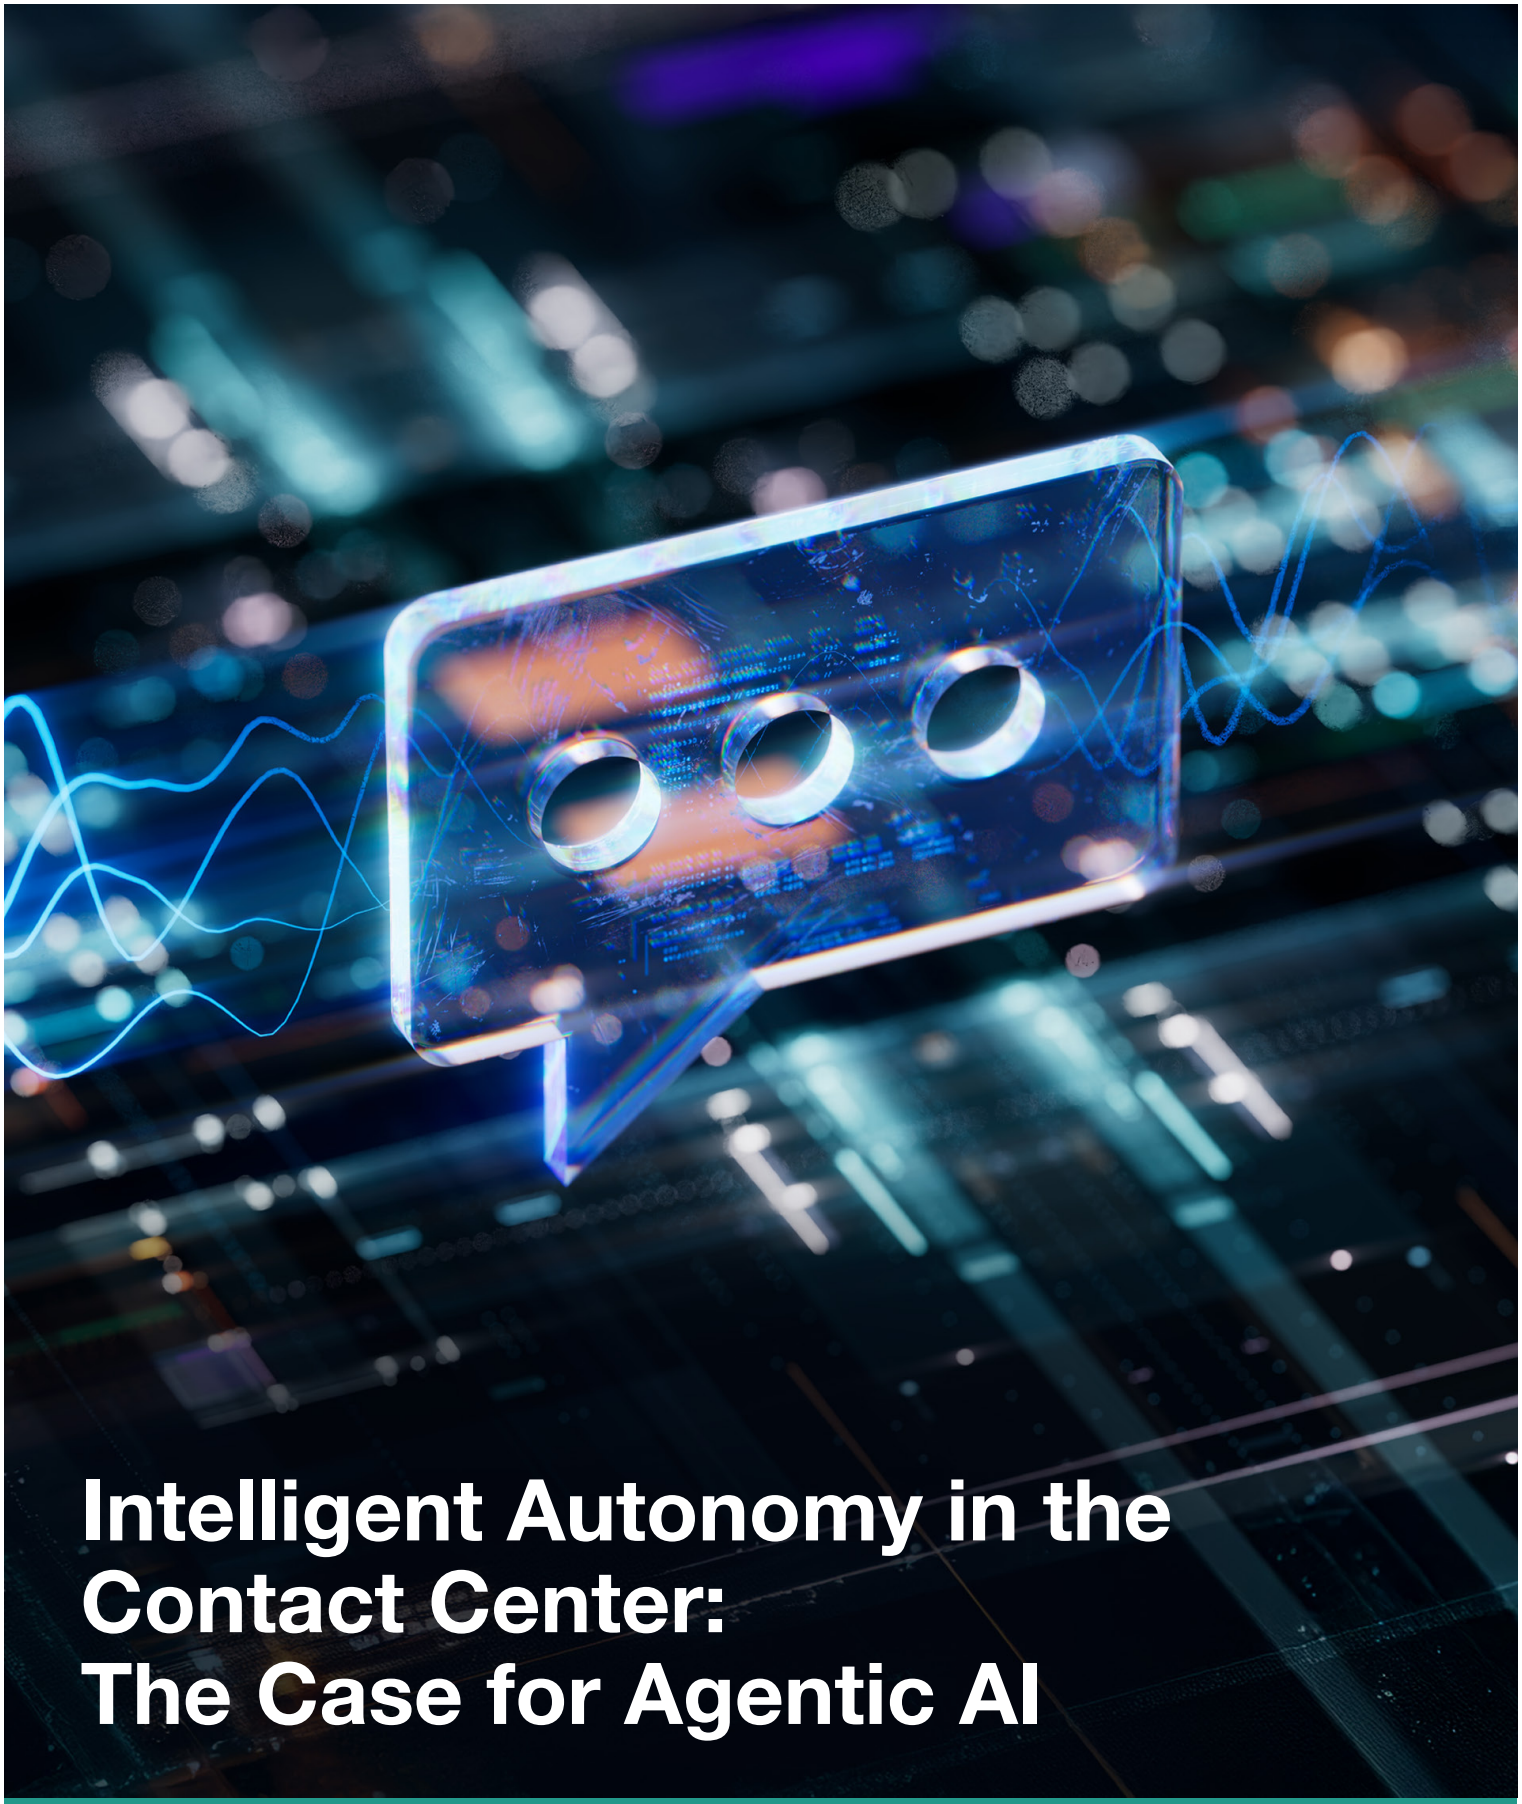

# Intelligent Autonomy in the Contact Center: The Case for Agentic AI

**Whitepaper by**

Adam Geffner, Head of AI and Global Contact Center Transformations, Mphasis

Mphasis

The Next Applied

# Contents

# 1. The Shift from Reaction to Initiative: Welcome to the Agentic Era

Does this sound familiar? A customer spots a suspicious charge on their bill and reaches out to custom support. They start interacting with a chat bot with limited ability to help. Then they call to speak to a live agent, wait on hold, and after re-verifying who they are, wait on hold some more. Then the agent begins to research the issue, or they have to transfer or call back to another dept, where the cycle starts all over again.

But what if they didn't need to go through all that? What if an intelligent system had already detected the anomaly, performed triage or root cause analysis and initiated resolution, adjusted their bill, and sent them a text or email about the issue and resolution proactively, before the customer ever reached out – all without any human involvement needed?

That's not a vision of the distant future. That's the emerging reality of Agentic AI and it's already in motion.

As organizations move beyond scripted bots, keyword-driven automation, and task-specific virtual agents, a new type of AI is taking shape — one that doesn't just wait for instructions but acts with intent.

**Agentic AI** refers to systems capable of pursuing goals, making autonomous decisions, and coordinating actions across tools, systems, and channels. These agents operate more like collaborators than tools: they plan, learn, adapt, and execute. They transform how work gets done in the contact center.

This shift isn't about replacing people. It's about elevating them, freeing human agents from repetitive tasks and enabling them to focus on empathy, nuance, and judgment. It's about moving from **reactive service** to **proactive engagement,** and from **scripted interactions** to **autonomous resolution**.

In this paper, we unpack what Agentic AI truly is (and isn't), not just as a concept, but as an operational model. We'll explore how enterprise-grade autonomy is being built, what separates Agentic AI from prompt-based LLMs, and how organizations are deploying it to rewire how contact centers function, a whole new level of autonomy and automation never achievable before. As with most new and exciting tech especially in AI, it's not without some risk, and we'll look at risks and concerns, primarily around hallucinations and accuracy.

Coined by Andrew Ng in June of 2024, "Agentic AI" is still evolving in meaning. What makes an AI Agentic? Where and how can it be applied and the ethical considerations around humans taking their hands off the proverbial steering wheel and where the boundary lies between traditional automation and true agency are still heavily debated. To understand what Agentic AI is, we must first understand what it means to have **agency**.

For people, it means to possess the **capacity to act independently**, make your own **choices**, and exert **influence or control** over your environment or circumstances.

> The best way to predict the future is to create it.
> **– Peter Drucker**

When a person has agency, they are responsible for deciding what actions to take. They initiate actions and behaviors rather than reacting or following a request. They own the consequences of their decisions and take other actions as a result of those consequences. Having agency is what allows individuals to shape their environment rather than simply respond to it. Agency is not just an action; it's intentional action, driven by goals. When we extend this concept to AI, we begin to see what makes agentic systems different.

Agentic AI refers to systems that have agency, the ability to act autonomously, initiate actions, make context-driven decisions pursuing goals and execute tasks across multiple steps to achieve larger goals and initiatives. Most importantly, to take on these tasks with minimal or no human intervention.

**So why is there still so much confusion about what qualifies as Agentic AI?**
Much of the uncertainty stems from differing interpretations of **how much autonomy** a system must exhibit to be considered truly agentic. This ambiguity is reminiscent of the early days of autonomous vehicles, when the industry lacked a commonly accepted model to distinguish between driver-assist features and full autonomy. AI faces a similar challenge. While the concept of agency is becoming clearer, the levels of initiative, planning, and decision-making vary widely between implementations. A **five-level maturity framework** has been created to address this thought, and it is gaining support. This model provides a structured way to classify Agentic AI systems by their capabilities, impact, and degree of autonomy.

| Level | Name | Description | Capabilities & Characteristics | Contact Center Impact | Example Use Case |
|---|---|---|---|---|---|
| Level 1 | Reactive Automation | Rule-based bots with no learning or memory | Stateless, rule-based, no context or adaptation | Handles FAQs, basic IVR; high human dependency | FAQ bots, IVR menus, scripted chatbots |
| Level 2 | Guided Execution | AI suggests action; human executes | Context-aware, single-task logic, human-led | Agent assist; reduces handling time | Suggested replies, Siri, Alexa |
| Level 3 | Supervised Autonomy | AI performs tasks with oversight | Multi-step reasoning, tool use, limited autonomy | Bot assist; automates routine tasks | Form fills,CRM updates, LangChain agents |
| Level 4 | Autonomous Agents | AI plans and acts independently | Goal-driven, memory, self-monitoring, API integration | End-to-end issue resolution; reduces FTE | Billing resolution Devin, DeepSearch agents |
| Level 5 | Collaborative Intelligence | AI and humans collaborate dynamically | Real-time learning, predictive, ethical, HITL | Proactive CX; real-time coaching and routing | Churn prediction CEO AI, researcher AI |

# 2. From Scripts to Autonomy: Why Agentic AI Matters Now

For years, contact centers have operated within the constraints of rule-based IVRs, scripted interactions, and bots that were useful – when handling common tasks, or more to the point, when questions matched the scripts.
The first bots were nothing more than visual IVR, and the more recent conversational AI bots have had mixed levels of success based on how well customer queries were aligned to the predefined interaction flows built.
These systems - while sufficient for high-volume, low-complexity needs - were never intended to manage nuance, intent, or the growing demand for personalized, real-time service.

From the first utterances of IVR systems in the early 70s to the more intelligent conversational AI systems over the past few years, there has nearly always been a discrepancy or delta between what customers were promised, and the real outcomes of what was delivered. Enter Agentic AI.

Agentic AI represents a leap forward. Unlike traditional automation that simply responds on cue, Agentic systems detect goals, plan actions, and execute complex workflows independently. It's not just reactive, it's proactive.
It doesn't follow rules; it reasons through them - initiating, adapting, and collaborating across channels and systems.

This shift is more than technical, it's strategic; and it's being driven by real pressure across the industry.

- *Rising customer expectations*
- *Shrinking operational margins*
- *Workforce constraints*
- *An explosion of systems and data sources that need coordination, not just access*

Agentic AI enables a new era of customer engagement, where automation becomes more capable and more contextually intelligent.

We've reached a point where rule-based tools can't scale the kind of experiences customers demand.

Unlike conventional bots, or even modern LLM-based assistants that rely on prompt-response logic, Agentic AI systems take initiative. They learn over time, manage multi-step tasks, and adapt to real-world conditions. In the contact center, that means AI no longer waits to be told what to do. It detects intent, evaluates context, and executes on behalf of the customer or the agent.

This isn't a minor upgrade. It's a foundational shift - from automation as a **tool** that waits to be told what to do, to automation as a **teammate** that understands what needs to be done and takes action. As we enter this new phase, Agentic AI stands not at the edge of possibility, but at the center of how customer experience will be defined in the coming years.

# 3. Catalysts Accelerating the Rise of Agentic AI

So, what's changed? Why is Agentic AI becoming not only possible, but essential? It's not due to a single breakthrough or even focused improvement in one area - it's a convergence. Critical shifts across technology, operations, and customer behavior have aligned to make Agentic AI viable at enterprise scale.

Five major catalysts have been fueling that rise:

- **The Maturation of Foundation Models:** Advances in LLMs like GPT-4, Claude, and Gemini have moved AI beyond pattern mimicry, chunking, and basic intent classification. These models have evolved into what are better described as Language Reasoning Models (LRMs) - capable of contextual reasoning, short-term memory persistence, dynamic intent detection, and increasingly fluent dialogue. Together, these capabilities provide the cognitive foundation that agentic systems build upon to support autonomous decision-making.
- **Orchestration that Bridges Silos:** Until recently, most AI was limited to the front end - classifying utterances, assisting agents, or managing simple flows. Others were siloed within narrow domains like performance monitoring, WFO, or WFM. True execution across systems happened elsewhere. That's changed. Modern orchestration layers now enable intelligent agents to operate across real enterprise environments—not just simulate conversation. By bridging APIs, databases, CRMs, and workflows, these systems allow agents to move beyond insight into execution—interacting with disparate systems, triggering real actions, and supporting autonomous decision-making at scale.
- **The Experience Gap is Growing:** Customers no longer tolerate "good enough" automation. They've grown increasingly frustrated with the results of digital leakage, and while conversational bots have improved, customers expect intelligent, real-time, personalized support - whether they're asking a question, disputing a charge, or changing their policy. Agentic AI is designed to meet this expectation, acting in real time with context and initiative.
- **Pressure to Scale Intelligently:** Contact centers face ongoing pressure to improve efficiency without increasing headcount. Agentic AI enables intelligent scaling - handling more volume without increasing proportional headcount. It relies more on what self-service can handle, offloading simpler interactions from agents, and in some cases, even eliminates the need for customer outreach altogether—by resolving issues proactively and without escalation.
- **Global Complexity Requires Smarter Automation:** Operating across geographies, languages, and time zones demands adaptable systems. Agentic AI's multilingual, cross-cultural capabilities make it uniquely positioned to handle this complexity at scale.

# 4. Agentic AI vs. LLM-only Systems: The Critical Difference

Large Language Models (LLMs) have redefined how machines interpret and generate human language. But language fluency alone doesn't equate to meaningful autonomy. Sustained context, real-world execution, and measurable outcomes are still required.

LLMs can converse, summarize, and perform basic reasoning within a single thread. But they stop short of acting independently and performing real-world actions. Think of them as a coach on the sidelines. Highly knowledgeable, able to analyze and advise, but not the one taking the shot or executing the next move. They're not out there playing the game. That's where **Agentic AI** steps in, building on LLM foundations, but extended with persistent, stateful memory, tool usage, and autonomous execution. It's a reversal of the "student becomes the master" analogy. Here, the coach becomes the player - entering the field, making real-time decisions, taking the shot, actively participating and driving results.

| Capability | LLM-only Bots | Agentic AI Agents |
|---|---|---|
| Language Fluency | ☑ Fluent | ☑ Fluent |
| Multi-turn Context Retention | ⦸ Limited | ☑ Persistent |
| Autonomous Execution | ⦸ Not supported | ☑ Supported |
| API/Tool Invocation | ⚠ Limited plugins | ☑ Robust integrations |
| Intent + Outcome Alignment | ⦸ Weak | ☑ Strong |

## 4.1 Limitations of LLM-only Systems

- Fluent in conversation but lack continuity across sessions
- Limited reasoning across multi-step or long-term tasks
- Unable to act independently across channels, tools, or workflows
- Reliant on humans to interpret, initiate, trigger, or finalize actions

## 4.2 Strengths of Agentic AI

- Built on LLMs but extended with orchestration logic, planning, and memory
- Maintain stateful memory and context for continuity and long-term personalization
  - A critical requirement for hyper-personalized services, which are just beginning to emerge
- Prioritize and execute complex, real-world workflows across CRMs, APIs, and other applications without HITL supervision
- Proactively deliver outcomes, not just responses, and at scale

**Bottom Line**: LLMs help machines understand; Agentic AI enables them to solve and act. It's the difference between answering questions and solving problems.

# 5. The Mphasis Approach: Operationalizing Agentic AI

Agentic AI shouldn't just respond - it should reason, act, and adapt, while still operating within frameworks that uphold control, trust, and governance.

At Mphasis, this philosophy aligns closely with our **Straight Through Processing (STP)** methodology, where intelligent automation flows seamlessly across the full customer journey - reducing friction, eliminating manual handoffs, and accelerating meaningful outcomes.

## 5.1 Balanced AI with Human-in-the-Loop (HITL)

Autonomy doesn't mean an absence of oversight. We advocate for a **Balanced AI Model**, where Agentic systems are empowered to act independently, but humans remain in the loop at critical decision points. This ensures responsible deployment; especially vital in high-stakes or sensitive interactions, where ethics, empathy, and regulatory alignment matter the most.

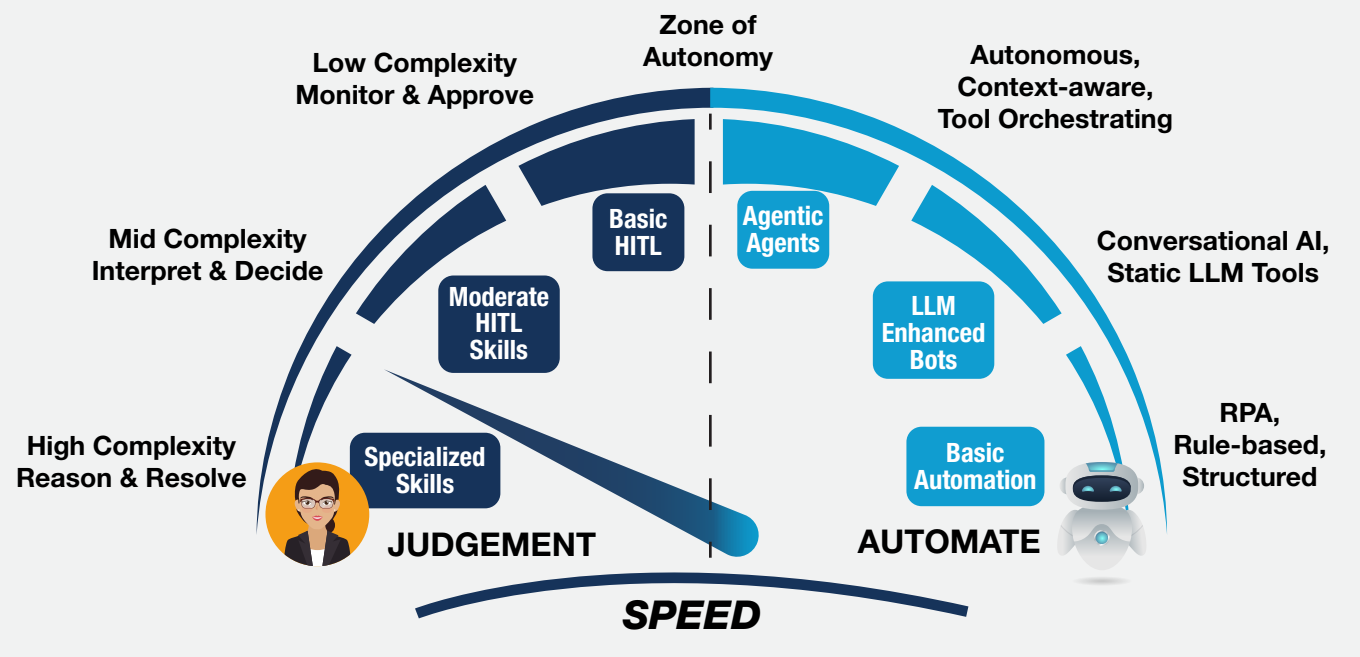


## 5.2 Core Pillars of the Mphasis Agentic AI Strategy

- **Purpose-built Agentic Frameworks**: Custom-designed agents capable of managing complex, multi-step service journeys—across voice, chat, email, and digital channels.
- **Real-time Orchestration Engines**: AI agents that dynamically coordinate APIs, customer profiles, enterprise data, and session history context to drive timely, context-aware actions in real time.
- **Advanced Intent & NLU Models**: High-accuracy understanding of language, goals, and context—fueling personalized interactions, intelligent redirection, and fast resolution.
- **Strategic Platform Ecosystem Partnerships**: Collaborations with AWS, Azure, Google Cloud, Salesforce, Kore.ai, and CCaaS leaders ensure flexibility and enterprise-grade scalability.

Our approach empowers enterprises to deploy Agentic AI architectures at scale and operationalize them with confidence, combining innovation with governance, accelerating outcomes while maintaining oversight, and balancing autonomy with accountability. This is how we achieve a balance of intelligent automation with human accountability.

# 6. Real-world Case Study: Tier-1 Banking CX Transformation

## 6.1 The Challenge

A global tier-1 bank was facing mounting pressure across its customer support operations:

- High call volumes were overwhelming human agents
- Long wait times were eroding customer satisfaction, leading to customer frustration
- Fragmented multilingual support was failing to meet expectations across global markets

## 6.2 The Solution

Mphasis deployed its Agentic AI platform to automate tier-1 service requests across both voice and digital channels. Deep integration with the bank's back-end systems enabled autonomous agents to manage entire workflows—across languages, channels, interaction types and customer intents – all without the need for human intervention.

## 6.3 The Result

The impact was both immediate and measurable:

- **20+ Languages Supported**: Delivering consistent, always-on global engagement
- **40% Reduction in Escalations**: Agentic AI resolved a significant share of interactions without human intervention
- **35% Decrease in AHT**: Faster resolutions improved operational efficiency, CX, and EX
- **Significant CSAT & NPS Uplift**: Customers experienced faster, more intelligent service, agents shifted focus on higher-value, advisory conversations

This deployment illustrates how Agentic AI doesn't just automate tasks, it transforms both the **Customer Experience (CX)** and **Employee Experience (EX)**, resulting in a better **Unified Experience (UX)** - delivering tangible results across operational, experiential, and strategic dimensions.

# 7. Future Outlook: Agentic AI and the Intelligent Enterprise

The contact center of the future will be unrecognizable from its legacy roots. No longer reactive service hubs, they are evolving into intelligent, AI-first ecosystems, where Agentic systems and human advisors collaborate seamlessly in real time to deliver intelligent, adaptive, and empathetic customer experiences.

The following trends signal how Agentic AI is set to reshape enterprise operations:

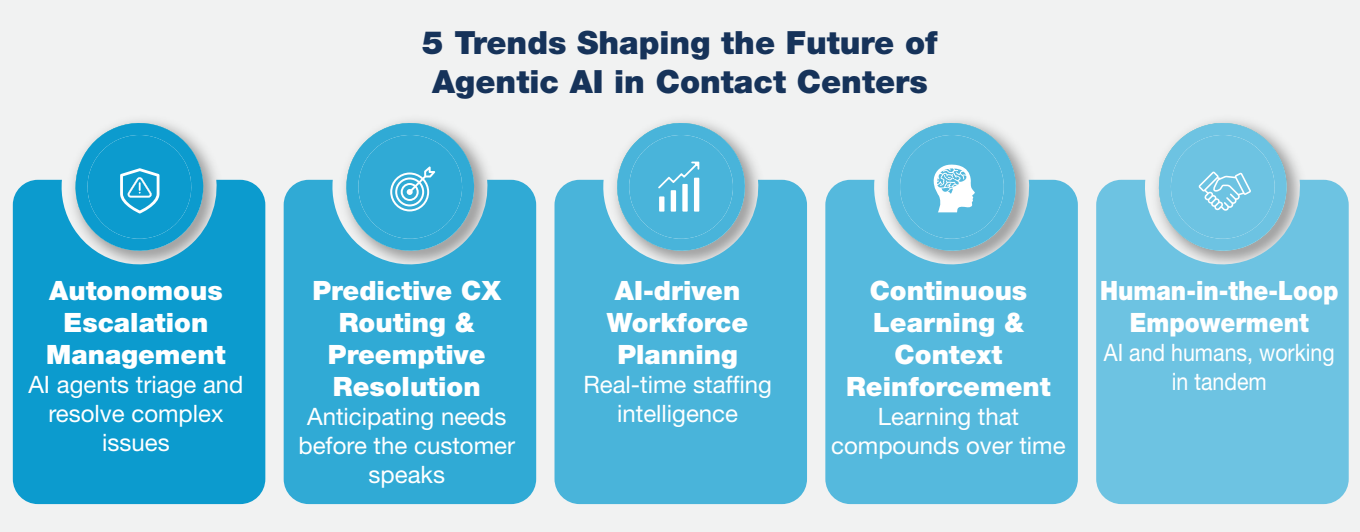

- **Autonomous Escalation Management**: AI agents will increasingly manage complex inquiries end-to-end—triaging issues, navigating decision trees, and resolving cases with minimal human involvement.
- **Predictive CX Routing and Preemptive Resolution**: Armed with behavioral data, actionable insights and real-time context, agentic systems will anticipate customer needs before they're expressed—deflecting known issues, surfacing proactive resolutions, and routing based on urgency, value, and intent.
- **AI-driven Workforce Planning (WFO)**: From forecasting volume spikes to optimizing agent skill allocation, intelligent systems will refine workforce strategies in real time—improving both staffing efficiency and service quality.
- **Continuous Learning and Context Reinforcement**: Agentic AI will not only learn from each interaction. They will retain and build upon context, creating compounding intelligence. Think of it as a self-evolving CI/CD system which iteratively improves intent and understanding – offering agents smarter recommendations and customers more seamless, personalized journeys.
- **Human-in-the-Loop (HITL) Empowerment**: As automation expands, human agents will play a vital role in governance, ethical alignment, and complex decision-making. They will focus more on oversight, ethical judgement and complex decisions. A hybrid model where automation enhances, rather than replaces, human contribution.

# 8. Risks and Considerations: Deploying Agentic AI Responsibly

Agentic AI holds transformative potential, but responsible deployment demands clear guardrails. Enterprises must anticipate and mitigate the following risks to ensure secure, ethical, and effective implementation.

## 8.1 Hallucination and Accuracy Risks

Not all AI outputs are grounded in fact. Foundation models trained on broad internet datasets—such as GPT-4, Claude, or Gemini—can occasionally generate speculative or incorrect responses. Understanding model behavior is essential, but also becoming more difficult as the models themselves continue to advance.

- **Open-Ended Models**: Powerful but prone to hallucination due to generalized training data.
- **Closed-Domain Models**: Training on curated, enterprise-specific datasets significantly reduces hallucination, improving reliability for business-critical use cases.

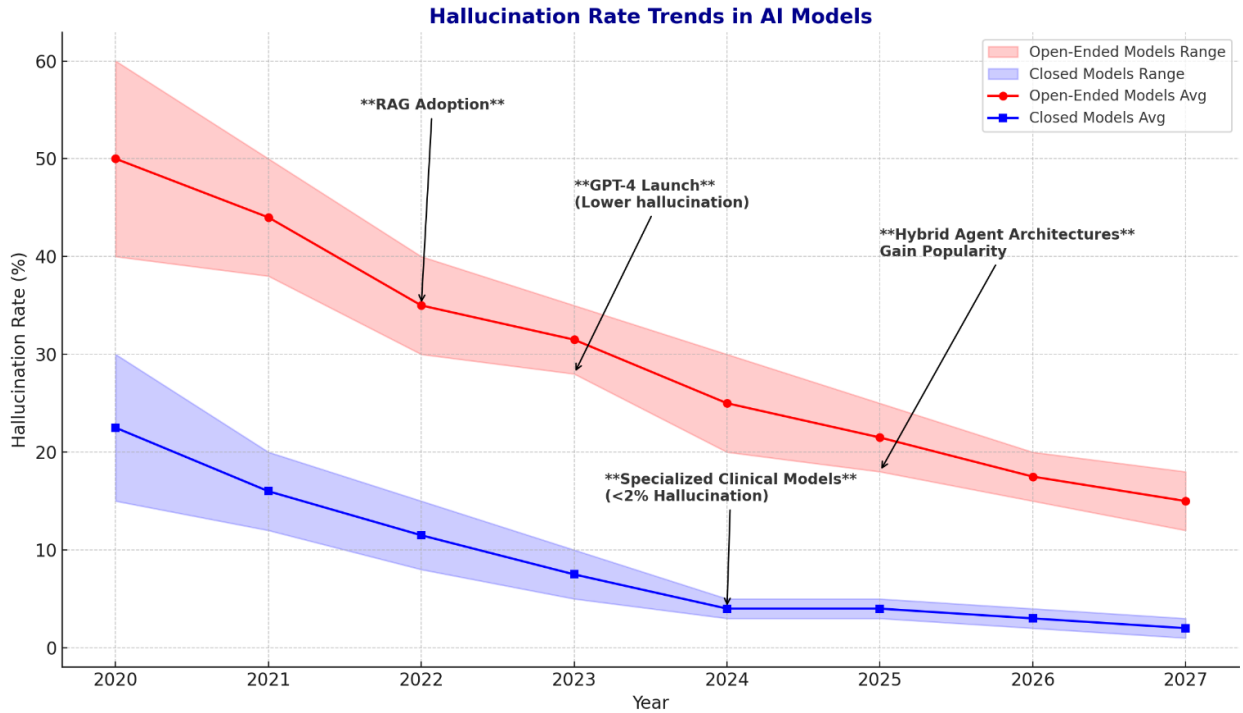

## 8.2 Explainability and Transparency

In highly regulated industries, decisions can't just be correct - they must be defensible. Agentic AI must offer:

- **Auditable Reasoning**: Systems should identify and log traceable logic paths and decision trees.
- **Regulatory Alignment**: Transparency that supports operational compliance and reduces violation risks.

## 8.3 Ethical Autonomy and Governance

As autonomy increases, so does the responsibility to ensure AI acts within boundaries. Governance mechanisms must evolve in tandem.

- **Human-in-the-Loop (HITL) Design**: Critical decisions remain under or validated by human supervision.
- **Governance Frameworks**: Policies that establish operational boundaries, acceptable behaviors, escalation paths, rollback triggers, API-level guardrails and fail-safes.
- **Bias Monitoring & Mitigation**: Ongoing monitoring to reduce algorithmic bias and ensure unbiased behaviors, especially in customer-facing or credit-related applications.

| Risk | Description | Why it Matters |
|---|---|---|
| Over-Autonomy/ Loss of Control | Agents acting beyond intended scope, especially when they have API/write-level access to enterprise systems | Even well-trained agents can cascade unintended consequences across systems without clear execution boundaries |
| Tool Misuse or Tool Fatigue | Agentic AI that invokes tools inappropriately, too often, or inefficiently (e.g., hammering APIs) | Can lead to degraded performance, cost overruns, or even service disruptions. Especially critical in cloud-native systems with API rate limits. |
| Security & Surface Expansion | More autonomous agents = more endpoint integrations, and potential attack surfaces | Autonomous actions through APIs, file uploads, or form-fills increase exposure to external threats and require security hardening |
| Model Drift & Behavior Degradation | As context or data shifts over time, models may produce increasingly inconsistent or biased outcomes | Continuous monitoring and retraining strategies are needed to preserve long-term accuracy and alignment |
| Decision Auditing & Replayability Gaps | Lack of structured logs or traceability into why the agent made a specific decision | Without decision traceability, it's difficult to audit or defend choices in regulated settings. This goes beyond explainability—it's about structured observability. |

# 9. Final Word: Leading the Next CX Transformation

The convergence of advanced LLMs, mature orchestration frameworks, and rising customer expectations has propelled Agentic AI from concept to enterprise reality.

This isn't just about automating interactions - it's about reimagining the contact center as an **AI-first, human-augmented ecosystem**. Agentic systems don't just understand intent, they act on it. They plan, adapt, and execute across tools, channels, and workflows to deliver real outcomes at scale.

This transformation is no longer optional. Customers now expect intelligence, immediacy, and empathy, not just scripts and queues. Enterprises must respond with automation that's both **resilient and responsible**, balancing innovation with oversight, autonomy with accountability.

At Mphasis, we bring this vision to life through purpose-built Agentic frameworks, real-time orchestration engines, and a deep ecosystem of platform partners. The result: measurable gains in efficiency, FCR, AHT, and customer satisfaction - without compromising compliance, trust, or human judgment.

The contact center of the future isn't an AI-versus-human battleground. It's a **symbiotic model** where intelligent agents elevate human capabilities and deliver smarter, more empathetic customer experiences at scale.

Those who embrace Agentic AI today won't just keep pace with change - they'll define what customer experience means for the next decade. The question isn't if Agentic AI will transform service. It's ***who will lead the charge***.

## About Mphasis

Mphasis' purpose is to be the "*Driver in the Driverless Car*" for Global Enterprises by applying next-generation design, architecture and engineering services, to deliver scalable and sustainable software and technology solutions. Customer centricity is foundational to Mphasis, and is reflected in the Mphasis' Front2Back™ Transformation approach. Front2Back™ uses the exponential power of cloud and cognitive to provide hyper-personalized ($C = X2C^{2}_{TM} = 1$) digital experience to clients and their end customers. Mphasis' Service Transformation approach helps 'shrink the core' through the application of digital technologies across legacy environments within an enterprise, enabling businesses to stay ahead in a changing world. Mphasis' core reference architectures and tools, speed and innovation with domain expertise and specialization, combined with an integrated sustainability and purpose-led approach across its operations and solutions are key to building strong relationships with marquee clients. Click here to know more. (BSE: 526299; NSE: MPHASIS)

For more information, contact: marketinginfo.m@mphasis.com

**USA**
Mphasis Corporation
41 Madison Avenue
35th Floor, New York
New York 10010, USA
Tel: +1 (212) 686 6655

**UK**
Mphasis UK Limited
1 Ropemaker Street, London
EC2Y 9HT, United Kingdom
T : +44 020 7153 1327

**INDIA**
Mphasis Limited
Bagmane World Technology Center
Marathahalli Ring Road
Doddanakundhi Village, Mahadevapura
Bangalore 560 048, India
Tel.: +91 80 3352 5000



WAS 05/08/25 US LETTER BASIL 9705

**www.mphasis.com**