HCLTech

Driving Agentic AI Adoption in Global Capability Centers (GCCs)

A repeatable, co-creation blueprint to scale from pilots to governed, production-grade autonomous execution

White Paper | February 2026

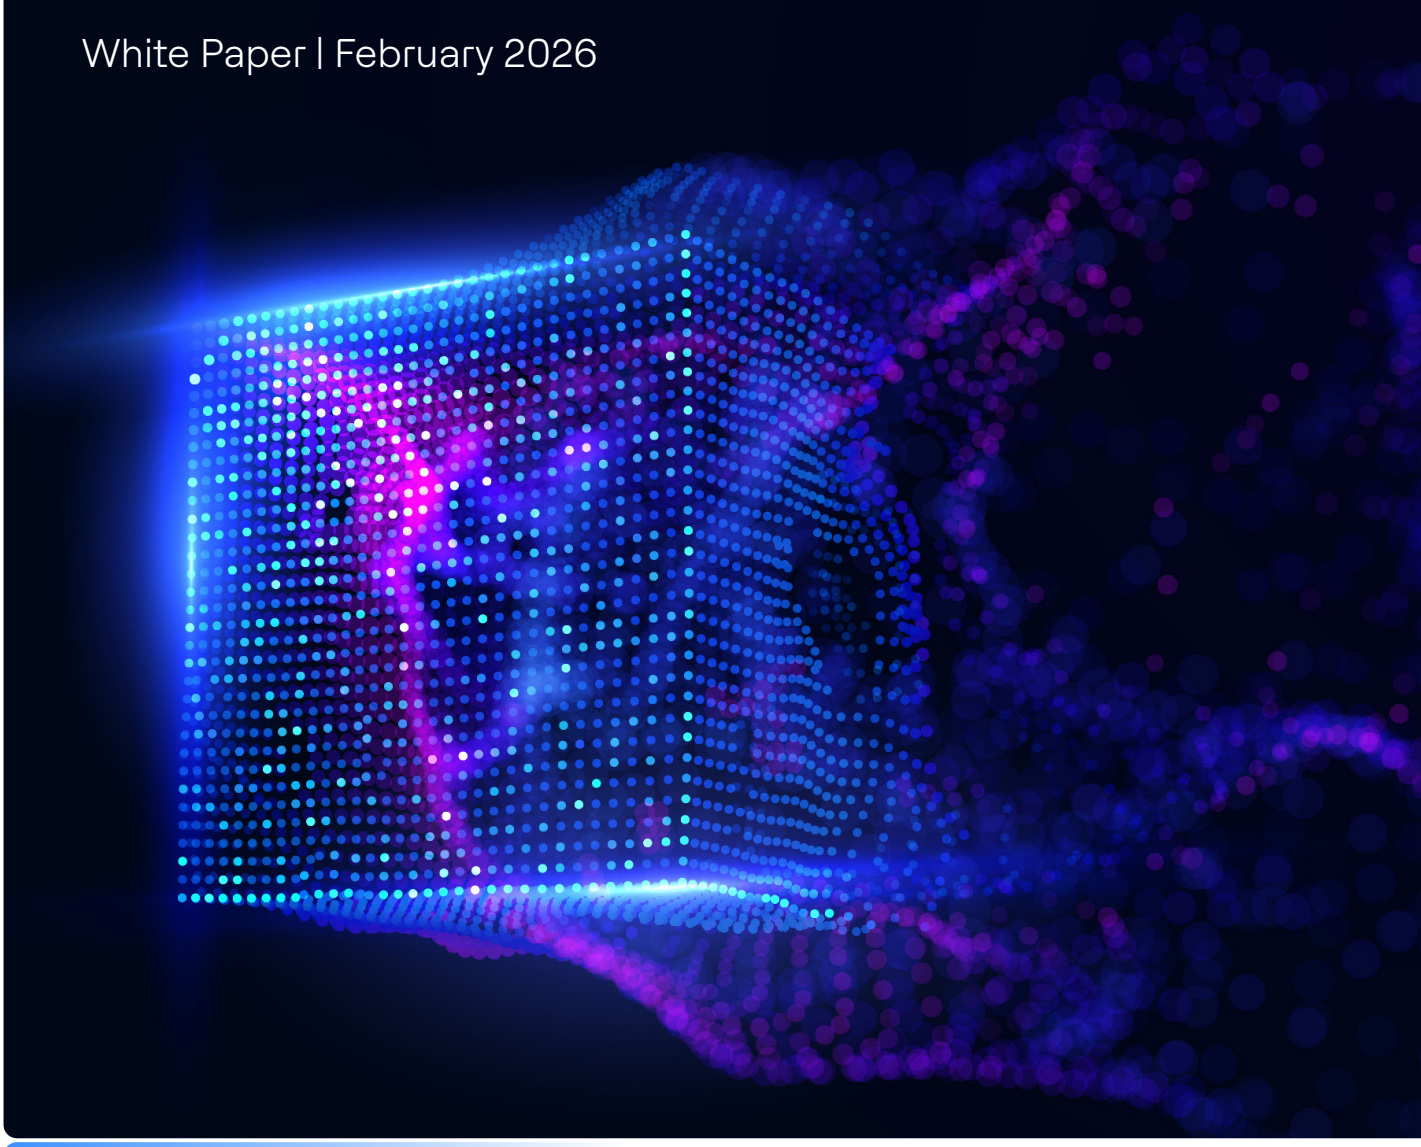

Table of Contents

Executive Summary

Enterprises are moving from digitization and automation into an era of cognition—where AI systems do not merely assist individuals, but begin to orchestrate enterprise workflows. Agentic AI marks this inflection point: goal-driven systems capable of multi-step reasoning, knowledge integration, tool use, and semi-autonomous or autonomous execution.

Global Capability Centers (GCCs) have evolved from cost-optimized delivery hubs into innovation engines responsible for modernization, AI adoption, and measurable enterprise outcomes. This structural position—domain proximity, data adjacency, and engineering density—makes GCCs natural control towers for scaling agentic capabilities across global enterprises.

Yet many initiatives remain stuck in pilot purgatory. The gap is not model potential; it is industrialization: legacy integration, data readiness, governance and security, platform orchestration, and the challenge of translating tacit domain knowledge into auditable agent behavior.

This white paper proposes a GCC-centric co-creation blueprint that combines People (a forward-deployed triad squad), Process (Value Stream Lifecycle Mapping), and Technology (HCLTech's AI platforms such as AI Force and AI Foundry) to move from PoCs to production-grade agent fleets with measurable business outcomes. This blueprint unifies the domain proximity of GCCs, the industrial engineering of service providers, and the deep research rigor of academic institutions like the Indian Institute of Technology (IIT) Kanpur. By anchoring research rigor with industrial pragmatism, academia provides the credibility and continuity necessary for large-scale adoption.

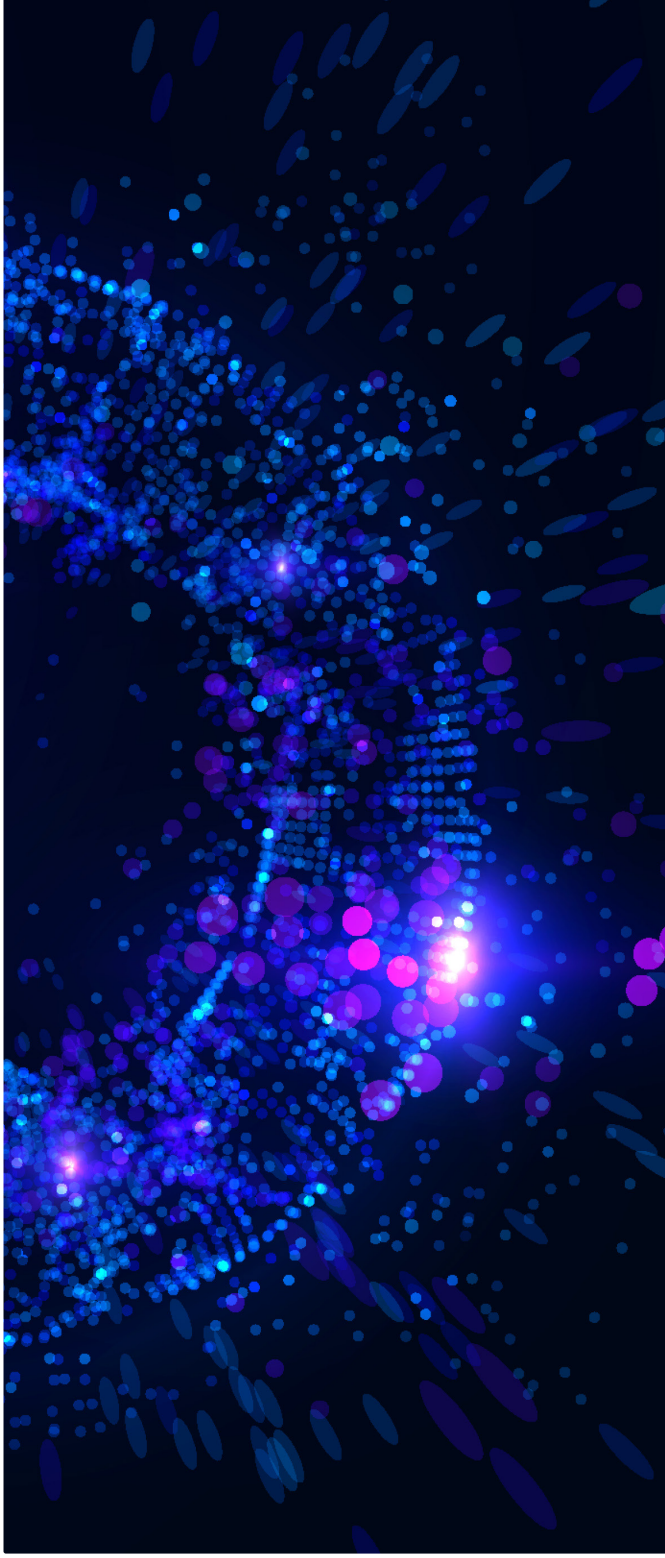

Abstract

This paper introduces a GCC–Academia–Service Provider co-creation framework for industrializing Agentic AI at enterprise scale. It unifies business domain ownership (GCCs), research rigor (academia), and industrial execution (service providers) into a closed-loop delivery model that supports scalable, governable, and economically feasible deployment of agentic systems across business domains.

The GCC Imperative: Why GCCs Are Central to Agentic Transformation

Global Capability Centers have evolved from cost arbitrage units into innovation powerhouses that sit at the intersection of global strategy and local execution. With over 1,600 GCCs in India alone, these centers now own critical business processes, technology platforms, and data assets that make them natural laboratories for agentic AI transformation. Unlike traditional IT departments constrained by legacy thinking or external consultants lacking operational context, GCCs possess the unique combination of business intimacy, technical capability, and organizational trust required to reimagine work fundamentally. Their cross-functional charter, proximity to domain workflows and data, and engineering prowess make them structurally positioned to become agentic control towers that design, deploy, and govern reusable agent frameworks across the enterprise.

Breaking Free from Linear Economics

The traditional GCC value proposition rests on linear economics: hiring smart people at competitive rates to execute defined tasks. This model faces inevitable compression as talent costs rise, attrition creates friction, and manual processes plateau in efficiency. Agentic AI enables GCCs to fundamentally disrupt this equation by enabling non-linear value creation. Instead of one analyst producing one report, an AI agent can simultaneously analyze hundreds of data sources, generate insights across multiple scenarios, and deliver personalized outputs at scale. The economics shift from headcount-based capacity to outcome-based velocity, where the same team of 100 can deliver work previously requiring 500, not through incremental automation but through wholesale process reimagination.

The Context Advantage for GCCs

Agentic AI implementation isn't just a technology deployment—it's a business transformation that requires deep contextual knowledge. GCCs hold this institutional memory of why processes exist, where exceptions hide, and how work actually flows versus how it's documented. This tacit knowledge proves critical when designing agents that must navigate ambiguity, make judgment calls, and integrate across fragmented systems. Technology vendors can provide tools and frameworks, but only GCC teams can answer questions like "Why do we manually verify these invoices but not those?" or "What implicit knowledge does our finance team use when they override the system?" This context becomes the training ground for agents that don't just automate tasks but augment decision-making.

While the structural advantages of GCCs make them the ideal control towers for AI transformation, the path from potential to production is rarely linear. The move from task automation to cognitive orchestration introduces a layer of complexity that current operational models are often ill-equipped to handle. This leads to a critical paradox: while GCCs are launching more pilots than ever before, they are simultaneously hitting a structural wall that prevents these innovations from reaching enterprise-grade scale.

The Deployment Chasm: Why many Agentic Programs in GCCs Stall Beyond PoCs

Across industries, GCCs have launched numerous AI initiatives, many of which demonstrate strong results in controlled environments—automating invoice processing with high accuracy, deploying HR chatbots, or building predictive customer churn dashboards. These early successes validate both the technology and the talent behind them. As organizations move from pilot to enterprise-wide deployment, however, the complexity of integrating with live systems, governance frameworks, and cross-functional workflows can slow momentum. What appears to be a seamless proof of concept must now operate within the full scale and variability of real-world enterprise environments. This dynamic is not a reflection of technological or capability limitations, but rather the natural transition from demonstrating innovation to industrializing solutions across complex operational ecosystems. With the right production frameworks, integration planning, and governance alignment, GCCs are well positioned to accelerate this journey from pilot to scaled impact.

01

The Absence of Industrial Grade Execution

Building a working AI model differs fundamentally from deploying it as a reliable business capability. GCCs bring strong process discipline and operational excellence to enterprise environments, providing a solid foundation for scaling AI initiatives. However, as AI adoption expands, productionizing solutions at scale requires additional specialized skills, adaptive engineering practices, and enhanced monitoring capabilities. With the right capability investments and governance frameworks, GCCs are well positioned to evolve into resilient AI production and lifecycle management hubs.

02

Missing Research Depth for Complex Problems

As GCCs evolve from automation initiatives toward more advanced, agentic AI capabilities, they begin addressing increasingly complex and dynamic problem spaces. These next-generation use cases often require expanded engineering approaches, adaptive architectures, and enhanced governance models to fully realize their potential at scale. How do you ensure an agent doesn't hallucinate when making financial recommendations? How do you build safety guardrails that block harmful actions without constraining legitimate flexibility? These questions require deep research expertise in areas like retrieval-augmented generation, multi-agent coordination, and AI safety that most GCCs don't maintain in-house. Without access to cutting-edge research, teams resort to brute-force prompt engineering, over-constrained rule sets that defeat the purpose of agentic flexibility.

03

Governance and Trust Deficit

Scaling AI from POC to production requires stakeholders to trust agents with real decisions that affect customers, compliance, and revenue. The challenge intensifies with agentic AI, where systems make autonomous decisions rather than just offering recommendations & insights humans can review. In the absence of robust frameworks for explainability, auditability, and progressive trust-building, AI initiatives can remain in extended pilot phases—prioritizing safety validation while the technology continues to evolve. Establishing strong governance models, transparent controls, and structured adoption pathways enables organizations to move confidently from experimentation to scaled deployment.

04

Change Management Challenges

Technology deployment fails most often not because the system doesn't work but because the organization doesn't change. Agentic AI threatens established workflows, decision rights, and performance metrics in ways that traditional automation doesn't. As AI and transformation initiatives scale, navigating these dimensions effectively often requires enhanced change management infrastructure, structured communication models, and enterprise-wide alignment. Strengthening these capabilities enables GCCs to drive adoption at scale while maintaining governance, stakeholder confidence, and operational continuity. Training programs focus on tool usage rather than mindset shifts. Incentive structures still reward individual output rather than human-agent collaboration. Without systematic change management aligned to the business context, even technically successful deployments face passive resistance that prevents adoption from spreading.

05

Integration Complexity

Enterprise environments consist of hundreds of interconnected systems, each with its own data formats, security protocols, and business logic. Agentic AI agents must navigate this complexity to deliver value. As AI initiatives progress, the integration effort can be more extensive than initially anticipated. Connecting intelligent agents to enterprise systems often involves substantial API development, data pipeline engineering, and security validation to ensure scalability, reliability, and compliance. Proactively planning for this integration complexity enables smoother deployment and faster time to value.

Closing this Deployment Chasm requires more than just better algorithms or faster infrastructure. It demands a fundamental shift in how people, processes, and technology are integrated into a single, repeatable delivery engine. To move ahead, organizations must adopt a co-creation framework that bridges the gap between deep research, industrial-grade engineering, and real-world domain expertise.

Crossing the Chasm: A Co-Creation Framework to Industrialize Agentic AI in GCCs

Closing this chasm requires a deliberate operating model—a repeatable blueprint that GCCs and their partners can apply across use cases, business units, and geographies.

This is the foundation of the GCC–Academia–Service Provider co-creation model. It is organized around three integrated elements—People, Process, and Technology—each designed to directly address one or more of the structural barriers that trap AI initiatives in the pilot stage.

People: The Forward-Deployed Triad ("Three-in-a-Box" Squad)

The most common reason AI projects stall between POC and production is that the people who understand the business problem, the people who can build the solution, and the people who can handle its technical complexity are all working in separate lanes—handing work off to each other rather than solving problems together. The co-creation model breaks this pattern by assembling a single, integrated squad where all three roles work side by side from day one.

GCC Business Domain Expert (Domain Owner): defines what success looks like, articulates the constraints and edge cases that only come from lived operational experience, and owns the change adoption journey within the organization. Crucially, they carry the tacit knowledge that no documentation captures.

HCLTech AI Implementation Expert: translates business intent into production-grade technical reality. This role designs the agent architecture, builds the integrations across enterprise systems, and handles the engineering disciplines—security, monitoring, failover, scalability, and operational readiness—that separate a working demo from a solution the business can rely on.

Academic Deep Research Expert: provides the technical depth for problems that go beyond standard software engineering. The academic partner—with deep expertise in areas like multi-agent systems, model evaluation, and AI safety—brings the research foundation to tackle these hard problems. Section 6 describes in detail how IIT Kanpur, our academic partner, plays this role in practice. Section 6 further highlights how IIT Kanpur, our Academic partner, would play a key role in driving this role forward.

Process: Value Stream Lifecycle Mapping

Agentic AI succeeds when implementation is guided by a disciplined, measurable process. The operating method proposed here is Value Stream Lifecycle Mapping—a seven-stage process that is based on value stream thinking (the practice of mapping exactly how work flows, where value is created, and where effort is wasted). It ensures that the squad is always working on the highest-impact opportunity and that each stage produces durable assets—reusable connectors, governance policy templates, evaluation frameworks, and operational runbooks—that make the next initiative faster and lower risk.

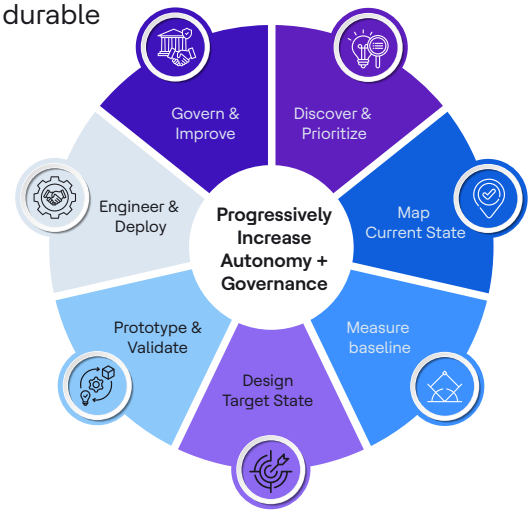

The intent at every stage is to answer three questions: Where do agents create the most value? How much autonomy is appropriate at this point in the journey? And what governance safeguards must be in place before expanding the system's reach? The table below maps each stage to the distinct contribution each role makes, making accountability clear and ensuring no expertise gap goes unaddressed.

Stage	GCC Domain Expert	Academic Research Expert	HCLTech Implementation Expert
Discover & Prioritize	Select value stream; define outcomes and constraints.	Assess feasibility; frame evaluation approach	Review platform fit; identify accelerators
Map Current State	Capture real workflow, exceptions, tacit knowledge.	Identify reasoning hotspots; propose knowledge representations.	Catalog systems, integrations, and security model
Measure Baselines	Validate KPIs; provide historical cases and metrics.	Define evaluation metrics and stress tests.	Set up telemetry and evaluation harness
Design Target State	Define autonomy tiers; approvals; exception paths.	Specify safety constraints and evaluation gates	Design orchestration, integration, and observability
Prototype & Validate	SME validation and feedback.	Run stress tests; drive robustness improvements	Build and iterate rapid prototypes
Engineer & Deploy	Lead change management, training, and runbooks	Identify drift risks; plan evaluation refresh	Harden for production; security and ops readiness
Govern & Improve	Manage KPIs, policies, and expansion roadmap	Research solutions for the hardest edge cases	Operate and optimize the live agent system

Technology: HCLTech Platforms as the Agentic Backbone

A co-creation model needs an execution backbone that reduces time-to-production. For GCCs attempting to scale agentic AI, the technology gap is often the most underestimated barrier — not because the right tools don't exist, but because assembling them coherently, integrating them with live enterprise systems, and operating them responsibly at scale requires an industrial-grade AI foundation that most GCCs cannot build from scratch.

HCLTech's full-stack AI ecosystem is designed to move enterprises beyond automation toward autonomous operations and adaptive value creation — at scale. It is architected around two strategic imperatives: **AI-Led Service Transformation and AI-Led Value Stream Innovation.**

Service transformation makes the GCC itself AI-native and operationally more effective. Value stream innovation then deploys that capability outward across enterprise-wide business processes. Together they create a complete technology arc — from internal delivery excellence to external business impact.

Strategic Imperative 1: AI-Led Service Transformation, Powered by AI Force

AI Force is HCLTech's GenAI and Agentic AI platform that automates and augments workflows across software and data engineering, IT operations, and enterprise business processes to significantly improve and accelerate outcomes. It is powered by enterprise-grade Responsible AI, auditability, and contextual intelligence — bringing together prompts, retrieval-augmented generation (RAG), tools, autonomous agents, governance, and multi-LLM flexibility in one place — enabling teams to build and deploy AI-driven workflows quickly and safely, complemented by three specialized workflow modules that streamline diverse enterprise operations. Each module targets a distinct GCC delivery domain:

AI Force.Software

Accelerates the software delivery lifecycle. It helps GCC engineering teams accelerate delivery and modernization with intelligent automation across the software lifecycle to improve quality and reduce time to market. For GCC teams managing large application portfolios, this means agents that generate code, run tests, flag incompatibilities during migrations, and modernize legacy systems — compressing timelines that previously required large specialized squads.

AI Force.Ops

Transforms IT operations from reactive to intelligent. It moves GCC operations teams from reactive to resilient, with automation and observability across hybrid cloud, digital workplace, network, cybersecurity, and application managed services to improve performance and experience. For GCCs running managed services at scale, this enables the shift from human-intensive ticket resolution toward autonomous detection, diagnosis, and remediation — the foundation of a NoOps model.

AI Force.Data

Makes data trustworthy and usable at speed. It turns complex data into insight with autonomous engineering and resilient data operations to help GCC data teams discover, trust, and use data faster. Given that data quality and availability are among the most common reasons agentic AI pilots fail in production, AI Force.Data addresses a root cause rather than a symptom — ensuring agents operate on reliable, well-governed information.

Strategic Imperative 2: AI-Led Value stream Innovation

While AI Force transforms how GCCs operate internally, Value stream innovation is where GCCs shift from executing predefined processes to orchestrating intelligent, adaptive business capabilities — the transition from service center to strategic innovation engine. This imperative is delivered through four distinct offerings.

AI Engineering

Under the AI Engineering suite, HCLTech offers two transformative capabilities designed to bridge the physical and digital worlds, enabling intelligent, adaptive, and scalable systems: Physical AI Solutions, focused on embedding intelligence into physical systems and environments; and Kinetic AI Solutions, centered on robotics and intelligent automation, enabling systems that move, interact, and adapt in real-world environments.

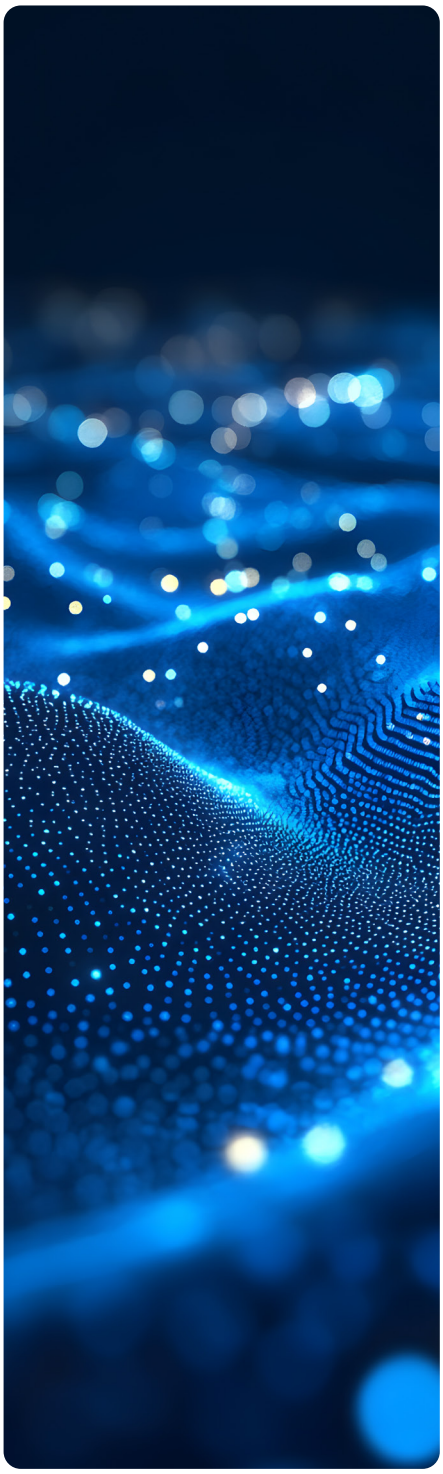

AI Factory

Provides the industrialized infrastructure backbone required to turn AI experiments into enterprise-grade deployments at scale. HCLTech AI Factory moves enterprise AI out of the lab and into the real world, making it practical and scalable. It brings people, processes, data, and platforms together inside a unified, governed, and industrialized ecosystem — addressing the platform complexity, fragmented data and models, disconnected tools, and limited governance that trap most AI initiatives in perpetual piloting.

Industry AI Solutions

Deliver the domain-specific intelligence that generic platforms cannot provide on their own. Industry AI Solutions are HCLTech's AI-led, IP-centric, scalable, and repeatable solutions designed to address sector-specific challenges. Built to deliver measurable business results, they enable enterprises to create value across key industry value chains.

AI Foundry

Serves as the enterprise-scale data and AI managed services layer that makes everything else sustainable. HCLTech AI Foundry is an enterprise-scale Data and AI managed services solution that helps organizations modernize data, simplify insights, and scale AI responsibly. AI Foundry addresses one of the most persistent failure modes in GCC AI programs: organizations that successfully build agents but lack the data infrastructure, model lifecycle management, and governance scaffolding to keep them performing reliably over

While the People-Process-Technology triad provides the operational backbone for deployment, the move toward fully autonomous agents is as much a scientific challenge as it is a technological one. For agents to navigate high-stakes enterprise environments without hallucination or failure, the implementation must be grounded in frontier R&D. This necessitates a strategic partnership with academia to provide the research rigor and safety guardrails required for production-grade AI.

The Intellectual Engine: The Foundational Role of Academia

The transition from AI as a tool to AI as an autonomous agent is a scientific transition as much as a technological one. Academic institutions provide the enabling conditions—including frontier R&D, a skilled workforce, and regulatory insights—that allow enterprises to bridge the gap between theoretical potential and production-grade execution.
Academia acts as a vital actor in the AI ecosystem, facilitating institutional interdependency between stakeholders such as governments and businesses. This role is not merely advisory; it is transformative, nurturing a culture where scientists and businesses co-create solutions to real-world problems rather than just co-existing. By meeting industrial scale with academic rigor, these institutions ensure that innovation is both locally grounded and cutting-edge. Coordinated efforts between academia and GCCs are essential for advancing national priorities, such as India's vision for "Viksit Bharat 2047," by developing high-value intellectual capital and ensuring technological self-reliance. Furthermore, the future of industrial engineering lies in this interdisciplinary synergy, where academic research in specialized domains—from medicine to aerospace—informs the design of intelligent and cyber-physical systems.

IIT Kanpur: The Strategic Research Pillar

Indian Institute of Technology Kanpur (IIT Kanpur) stands as a strategic research pillar in India's deep-tech ecosystem, uniquely positioned to power industrial-scale Agentic AI adoption. With decades of excellence in computer science, control systems, robotics, cybersecurity, and applied mathematics, the institute provides the intellectual engine required to design, validate, and deploy autonomous, decision-making AI systems. Its strong culture of interdisciplinary research enables the fusion of machine learning, systems engineering, and domain-specific expertise—an essential foundation for building reliable, production-grade agentic architectures rather than experimental prototypes.

Translating Research to Industrial-Scale Execution

Beyond academic excellence, IIT Kanpur offers the infrastructure backbone critical for scaling Agentic AI from lab to industry. Advanced computing facilities, secure research environments, robotics labs, and dedicated technology centers create a sandbox where autonomous systems can be stress-tested under real-world constraints. Centers such as its cybersecurity and AI research hubs contribute to developing resilient, safe, and explainable AI agents—key requirements for adoption in sectors like manufacturing, defense, finance, healthcare, and critical infrastructure. This infrastructure ensures that agentic systems are not only intelligent, but robust, compliant, and enterprise-ready.

Equally important is IIT Kanpur's strong industry interface and translational research model. Through startup incubation, industry-sponsored labs, and strategic partnerships, the institute bridges theoretical innovation with commercial deployment. It produces talent trained not just in algorithm development, but in systems thinking, optimization, embedded intelligence, and scalable architecture—skills essential for building autonomous agents that operate across complex enterprise environments. This steady pipeline of researchers, engineers, and entrepreneurs accelerates the transition from AI tools to fully autonomous decision ecosystems.

In positioning itself at the intersection of foundational research, high-performance infrastructure, and industry collaboration, IIT Kanpur functions as more than a university—it becomes a national strategic asset for Agentic AI. By anchoring research rigor with industrial pragmatism, it provides the credibility, capacity, and continuity necessary for large-scale adoption. As industries move from automation to autonomy, IIT Kanpur is poised to supply the intellectual leadership and systems-level innovation that define the next generation of AI-driven transformation.

The synthesis of these three pillars—GCC domain expertise, industrial-grade engineering from service providers, and the deep research rigor of academic institutions like IIT Kanpur—creates the necessary conditions for large-scale adoption. This collaborative model does more than just solve technical hurdles; it provides a roadmap for the fundamental organizational transformation required to lead in the age of autonomy.

Closing Summary: From Pilots to Autonomous Enterprise Advantage for GCCs

Agentic AI represents a fundamental shift—from automating tasks to orchestrating outcomes. For GCCs, this shift is not optional; it is the next stage of their evolution from delivery engines to strategic control towers of enterprise intelligence. As this paper has shown, the challenge is no longer about proving what AI can do, but about operationalizing what it should do—safely, reliably, and at scale.

The path forward lies in deliberate co-creation. By uniting GCCs' deep domain context, service providers' industrial-grade engineering, and academia's research rigor, enterprises can move decisively beyond pilots. The proposed People–Process–Technology blueprint provides a pragmatic, repeatable approach to crossing the deployment chasm—transforming fragmented experiments into governed fleets of autonomous agents that deliver measurable business outcomes.

Ultimately, successful Agentic AI adoption is as much an organizational transformation as it is a technological one. GCCs that embrace this model will not merely deploy smarter systems; they will redefine how work is designed, decisions are made, and value is created. Those that act now will shape the operating model of the autonomous enterprise—while others will be left optimizing yesterday's workflows in a world that has already moved on.

HCLTech


भारतीय प्रौद्योगिकी संस्थान कानपुर
INDIAN INSTITUTE OF TECHNOLOGY KANPUR