AI-FIRST GCC · POINT OF VIEW

The Next Frontier: AI-First GCCs

GCCs grew from cost centres into centres of excellence. The next step is a centre designed around AI from the start, and the gap between buying AI and governing it is where the value will be decided.

Vertara Global · Point of View · October 2026

AT A GLANCE

- India's 2,117 GCCs now generate $98.4 billion in revenue, and more than 1,200 have embedded AI and machine learning capabilities[1].
- Investment is near universal, with 83% of centres investing in generative AI, yet only one in five companies has a mature model for governing autonomous agents[2,3].
- An AI-first GCC redesigns work, talent and governance around AI. The first move is a governed workflow, not another tool licence.

01 The inflection point

GCCs have moved through two eras. The first was cost arbitrage: back-office execution and capacity added at lower cost. The second was the centre of excellence: global product ownership, specialised domain hubs and digital modernisation. Generative AI and agentic workflows now open a third, the AI-first GCC, in which AI is the operating foundation rather than a tool added to existing processes.

The difference is design. An AI-enabled centre gives people better tools. An AI-first centre changes what work people do, how teams are shaped and who is accountable when an agent makes a decision. Deloitte's 2026 survey of 3,235 leaders found 37% of organisations using AI at a surface level, with little or no change to existing processes[4]. That is the trap an AI-first GCC is built to avoid.

> *An AI-enabled centre adds tools to old workflows. An AI-first centre redesigns the workflow, and governs it, before it scales the team.*

02 Why 2026 is the turning point

Nasscom and Zinnov count 2,117 GCCs across 3,728 units in India, employing about 2.36 million people, with an AI talent base of 250,000. Nearly half the centres set up since FY2021 had AI as a core focus from launch, and

[1] Nasscom and Zinnov, "GCC Value Orbit: From Delivery Engine to Enterprise Nerve Centre", as reported by The Tribune (July 2026). www.tribuneindia.com/news/business/india-s-gccs-are-increasingly-leading-the-ai-mandate-for-global-enterprises-driving-global-value-creation-nasscom-zinnov-report/

[2] EY, GCC Pulse Survey 2025, "58% GCCs in India investing in agentic AI" (November 2025). www.ey.com/en_in/newsroom/2025/11/58-percent-gccs-in-india-investing-in-agentic-ai-two-third-creating-dedicated-innovation-teams-to-globalize-ideas-ey-gcc-pulse-survey-2025

[3] Deloitte, State of AI in the Enterprise 2026 (survey of 3,235 leaders, August to September 2025). www.deloitte.com/us/en/what-we-do/capabilities/applied-artificial-intelligence/content/state-of-ai-in-the-enterprise.html

[4] Deloitte, see note 3.

more than 250 dedicated AI centres of excellence now operate[5]. EY's 2025 GCC Pulse survey found 83% of centres investing in generative AI and 58% in agentic AI, with another 29% planning to scale it within a year[6].

Adoption is running ahead of operating change. In LinearB's 2026 benchmarks, even the top 10% of engineering organisations see AI assistance on 54% of pull requests but autonomous agents writing fewer than 5%[7]. Spend is not yet leverage, and the centres that close that gap first will set the terms for everyone else.

03 Looking ahead: 2027 and beyond

Three shifts follow for GCC leaders:

- **Mandates move from delivery to ownership.** Nasscom and Zinnov expect nearly 75% of GCCs could evolve into portfolio or transformation hubs within five years[8]. Ownership of enterprise AI capabilities will be part of that mandate.
- **The site leader becomes the AI governor.** 64% of site leaders already hold dual mandates covering areas such as AI governance[9]. Expect it to become an explicit, measured accountability.
- **Agents move from pilot to production.** Deloitte expects agentic AI use to rise sharply over the next two years[10]. Centres that build guardrails now will scale faster than those retrofitting them.

04 Four pillars of an AI-first GCC

A shared data and context core. Agents are only as good as the context they can reach. The GCC builds and maintains the enterprise knowledge layer, data pipelines and evaluation sets that let models work across finance, supply chain and engineering, and adapts models to proprietary data where it pays.

A talent model built on capability density. The key metric shifts from headcount to capability per person. Engineers move from writing every line to directing and reviewing agent work; domain specialists pair industry depth, such as actuarial or clinical knowledge, with practical AI fluency.

Governance as an operating discipline. Model evaluation, guardrails, audit trails, cost tracking and compliance with frameworks such as the EU AI Act sit with the centre, not with each project team. This is where the one-in-five maturity gap is closed or widened.

An ecosystem, not a fortress. More than 90% of leading GCCs work with universities and more than half co-innovate with startups[11]. Structured partnerships bring in talent and proven solutions faster than building everything in-house.

05 What changes by function

The shift looks different in each function, but the pattern is the same: an activity metric is replaced by an outcome metric.

[5] Nasscom and Zinnov, see note 1.

[6] EY, see note 2.

[7] LinearB, "The AI engineering productivity gap: 2026 benchmarks". linearb.io/resources/ai-engineering-productivity-gap

[8] Nasscom and Zinnov, see note 1.

[9] Nasscom and Zinnov, see note 1.

[10] Deloitte, see note 3.

[11] Nasscom and Zinnov, see note 1.

EXHIBIT 1 · FROM CENTRE OF EXCELLENCE TO AI-FIRST (ILLUSTRATIVE)

Function	GCC 2.0 baseline	AI-first GCC	What to measure
Software engineering	Distributed teams writing and testing code by hand	Agent-assisted coding, testing and maintenance behind review gates	Share of merged work written by agents; cycle time
Finance and accounting	Shared services running manual reconciliations	Continuous close with agent-assisted anomaly detection	Days to close; exceptions needing people
Customer operations	Scripted tier-1 and tier-2 support	Agents resolving multi-step requests, with human escalation	Resolution without hand-off; quality scores
Engineering R&D	Feature backlogs and legacy code	Simulation, digital twins and AI-assisted design	Time to prototype; reuse of validated models

Illustrative examples, not benchmarks. Targets should be set from each centre's own baseline. Vertara Global analysis.

That swap, more than any single tool, is what makes the transition visible to the parent. A board can see days to close fall; it cannot see licences activated turn into value.

06 How to run the transition

Run the transition in three phases, each ending in a gate the parent signs off before the next begins.

Phase	Focus	Decision gate
Months 1–6: Foundation	Audit data readiness; standardise toolchains and model access; roll out assistants with usage baselines	Are data and access fit for production agents?
Months 7–18: Workflow redesign	Rebuild two or three core processes around governed agents; stand up evaluation, cost and risk tracking	Measured gains on outcome metrics, within risk limits?
Months 19–36: Enterprise ownership	Re-anchor GCC KPIs to product velocity and enterprise outcomes; own shared AI capabilities	Does the parent treat the GCC as owner of the AI agenda?

- **Start with workflows, not tools.** Pick processes with clear outcome metrics and enough volume to show a result within two quarters.
- **Fund governance up front.** Budget evaluation, guardrails and cost tracking as a line item, not an overhead absorbed by projects.
- **Reskill before you restructure.** Move people into orchestration and review roles as agents take on volume work.
- **Report leverage, not adoption.** Licences and prompt counts tell a board little; agent-written work that ships and cycle-time gains tell it a lot.

QUESTIONS FOR THE BOARD

1. Which enterprise outcomes, not activities, should the GCC own in three years?
2. Which two or three workflows will we redesign first, and how will we measure the gain?
3. Who is accountable for AI governance at the centre, and with what authority?
4. How much of our AI spend is producing measured leverage today?

07 How Vertara Global helps

We start with a readiness diagnostic: the centre's mandate, data and toolchain maturity, workforce skills and governance, set against what an AI-first operating model would require. The output is a sequenced roadmap: the workflows to redesign first, the metrics to track and the governance model to put in place.

For new centres, we design AI into the operating model, entity, talent plan and workplace from day one. For existing centres, we help restructure toward AI-first ownership without disrupting delivery. Whichever operating model the client chooses, we stay engaged into year two, when the gap between adoption and leverage usually shows. The test of success: AI spend that shows up as measured enterprise outcomes.

ABOUT VERTARA global

Vertara Global is a practitioner-led GCC advisory firm. We help mid-market companies across sectors and geographies design, set up, scale and transform their Global Capability Centres, from the first business case and location choice to the operating model, entity, talent and workplace.

Our services span five pillars:

1. **GCC Consulting**
2. **Legal & Entity Support**
3. **People Advisory**
4. **Digital & AI**
5. **Engineering GCC Solutions**

Client-first and leader-led: senior practitioners run every engagement, and we stay involved after handover, because year two is where many GCCs are tested.