

The Agentic AI-first Global Capability Center

Architecting the autonomous
intelligence center

13 March 2026



The better the question.
The better the answer.
The better the world works.



Shape the future
with confidence

Contents

| | |
|---|----|
| Executive summary | 03 |
| 1. The structural evolution of GCCs | 04 |
| 2. The enterprise AI inflection point | 06 |
| 3. The agentic AI paradigm | 08 |
| 4. GCCs as enterprise AI launchpads | 10 |
| 5. Industry archetypes and impact profiles | 12 |
| 6. Case Illustration – AI-first Finance and Supply Chain GCC | 14 |
| 7. The four pillars of an AI-first GCC | 16 |
| 8. Governance, risk and responsible ai architecture | 18 |
| 9. Workforce transition and capability strategy | 20 |
| 10. Implementation roadmap and phased financial Model | 22 |
| 11. Strategic implications for enterprise leadership | 24 |
| 12. Conclusion: Architecting the autonomous intelligence center | 26 |
| 13. Authoring team | 28 |
| 14. EY GCC core team | 29 |
| 15. About AMCHAM | 30 |



Executive summary

Global Capability Centers are entering a decisive phase of structural reinvention. Over two decades, they have evolved from labor arbitrage engines into globally integrated delivery platforms and, more recently, into innovation hubs supporting enterprise digital transformation. Yet a more fundamental shift is now underway—one that demands a different conceptual lens and a different set of strategic choices.

The rapid maturation of generative AI, the emergence of agentic systems capable of autonomous multi-step orchestration, and the enterprise-wide pivot toward AI intelligent operating models are redefining what a GCC can—and must—become. Enterprises are no longer debating whether AI will influence their operating models. They are confronting how swiftly they can embed AI into their execution architecture while preserving governance integrity, operational resilience, and strategic control.

Incremental automation is insufficient. The competitive advantage accrues to those who redesign the activity architecture of the enterprise itself.

In this environment, retrofitting AI onto legacy process stacks yields modest, fragile gains. The real opportunity—and the real risk of inaction—lies at a different level: the redesign of the activity systems through which enterprises create and capture value. This is not merely a technology story. It is an operating model story.

This paper advances a clear and conservative thesis: the next generation of GCCs will not simply be digital-first. They will be Agentic-AI-first by design. An AI-first GCC embeds autonomous agents into its workflows from inception. It treats AI as a structural execution layer rather than a productivity tool. It re-architects processes around orchestration rather than hierarchy, redefines workforce models around augmentation and oversight rather than repetitive execution, and positions itself as a sovereign intelligence hub for the global enterprise. The economic case is compelling even under conservative modeling. Critically, the evidence base suggests that the path to this outcome is not linear. Different GCCs will arrive through different combinations of conditions—technology posture, digital ownership, innovation culture, GenAI readiness, operating model maturity, and transformation investment. Understanding these configurational pathways, not just the destination, is central to the strategic agenda facing GCC leaders today.

This is not automation. It is operating model reinvention.



CHAPTER 1

The structural evolution of GCCs

The GCC model has passed through four distinct maturity waves, each representing a structural shift in the mandate, operating logic, and value proposition of these centers. Understanding this trajectory is essential context for the case that follows.

GCC 1.0 – Cost Arbitrage

Early centers were designed to exploit global labor cost differentials. The mandate was operational efficiency and standardized service delivery. Metrics emphasized cost savings, headcount leverage, and throughput. Value was created by performing the same work, more cheaply.

GCC 2.0 – Global Business Services

As maturity increased, GCCs integrated multiple functions under unified governance. They assumed end-to-end process ownership, established Centers of Excellence, and began driving service quality alongside cost efficiency. The organizational unit became more complex, but the underlying logic remained execution-centric.

GCC 3.0 – Innovation Hubs

Over the past decade, leading GCCs embedded analytics, digital engineering, and transformation capabilities. They began influencing enterprise roadmaps and driving automation programs. Yet these models remained fundamentally process-centric. Traditional levers—Lean, ERP consolidation, shared services integration, and robotic automation—delivered incremental gains typically in the 5-15% range. They addressed repetitive tasks but struggled with cross-functional orchestration, exception-heavy workflows, and unstructured decision environments.

The Efficiency Plateau

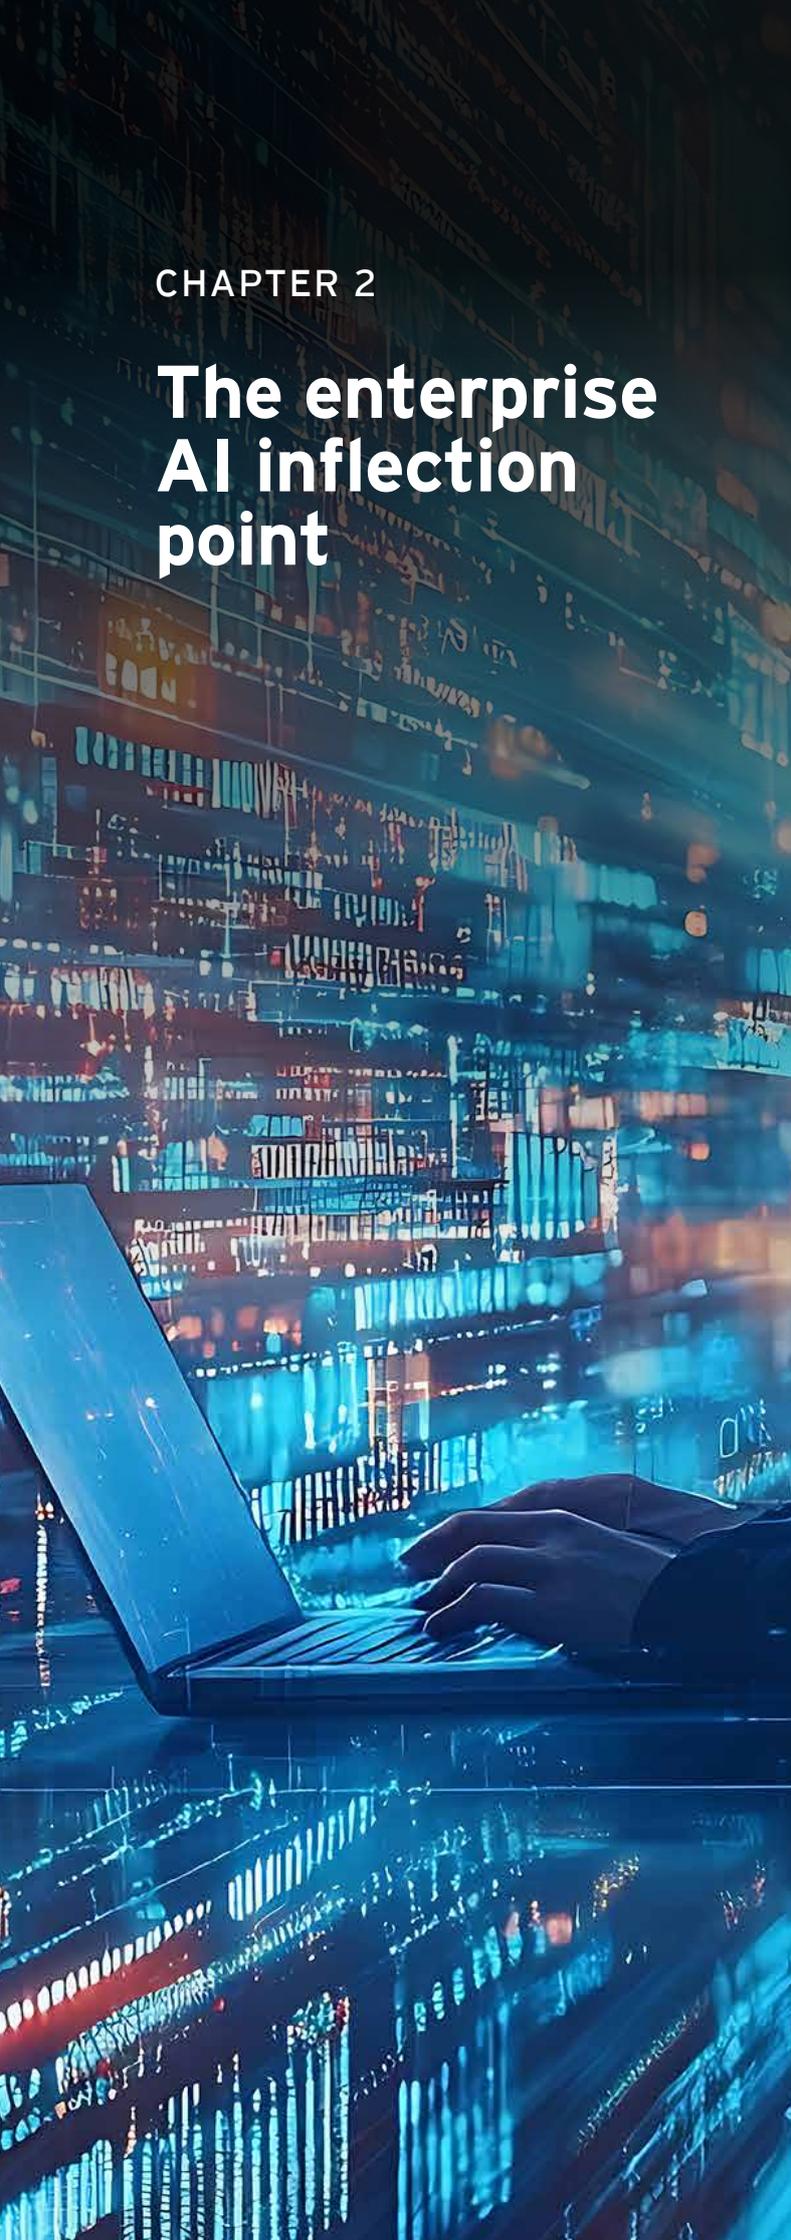
The tools that powered GCC 3.0 were built for a world of structured, predictable workflows. They are reaching their natural ceiling precisely as enterprise complexity—driven by data proliferation, ecosystem interdependence, and geopolitical volatility—is accelerating. The next wave of GCC value creation demands a different architecture.

GCC 4.0 – The Agentic-AI-first GCC

This emerging phase represents a structural break rather than an incremental advance. AI becomes the default execution layer. Workflows are redesigned around autonomous orchestration. Human effort becomes exception-based rather than process-based. Value shifts from incremental efficiency to autonomous value creation.

What distinguishes GCC 4.0 is not merely a new technology deployment. It is a redesign of the activity system itself—what the GCC does (content), how its activities are linked across functions and ecosystems (structure), and who controls the critical execution layer (governance). Centers that navigate this transition deliberately will establish structural differentiation that is difficult for competitors to replicate.





CHAPTER 2

The enterprise AI inflection point

Enterprise AI adoption has moved decisively beyond experimentation. Multiple large-scale surveys indicate that nearly half of major enterprises now have AI use cases in active production. However, most deployments remain functionally siloed—improving isolated tasks without redesigning the operating models in which those tasks are embedded.

~50%

of large enterprises have AI in production

yet most deployments remain task-level, not system-level¹

24% + 42%

of enterprise tasks are automatable or augmentable

implying two-thirds of task time is structurally redesignable²

Three macro forces are converging to shift the strategic calculus:

Falling AI compute costs, which have reduced the unit economics of inference by orders of magnitude, making large-scale deployment economically viable

Advances in reasoning-capable foundation models, which have extended AI applicability from structured data processing to complex, multi-step judgment tasks

Sources:

1. MIT NANDA: *The GenAI Divide - State of AI in Business 2025*

2. EY: *How much productivity can GenAI unlock in India? The Aldea of India: 2025*

Enterprise pressure for productivity and resilience, accelerated by post-pandemic cost scrutiny, talent market shifts, and geopolitical supply chain stress

As AI becomes more reliable and affordable, the constraint is shifting from technical feasibility to architectural courage. Enterprises that redesign their execution layers will capture structural advantage. Those that merely automate existing processes will realize incremental and rapidly commoditized gains.

The constraint is no longer what AI can do. It is whether the enterprise is willing to redesign the architecture in which AI operates.

GCCs are strategically positioned at this inflection point. They combine scale, governance maturity, cross-functional integration, and enterprise-wide data access. Critically, they offer something that distributed business unit deployments cannot: concentrated orchestration capability. Fragmented AI deployment—where each function builds independently—increases redundancy, governance complexity, and the risk of misalignment with enterprise architecture. GCCs, by contrast, offer the conditions for disciplined, scalable, and reusable AI deployment.

Recent data reinforces this momentum: 83% of GCCs in major markets are already engaged with generative AI solutions, while 58% are actively building agentic AI capabilities. The question for enterprise leadership is no longer whether GCCs will participate in the AI transformation. It is whether they will lead it by design, or follow it by default.



CHAPTER 3

The agentic AI paradigm



Agentic AI systems differ fundamentally from traditional automation—and from the first generation of generative AI tools deployed for task assistance. Rather than executing predefined rules or responding to isolated prompts, agentic systems perceive context, plan multi-step actions, invoke tools and APIs, coordinate across systems, and adapt dynamically to changing conditions.

In enterprise environments, this enables a qualitatively different category of capability:

Orchestration. End-to-end workflow orchestration, where agents manage process sequences from initiation to completion without human handoffs at each step

Exception resolution. Intelligent exception handling, where agents detect anomalies, assess context and resolve or escalate with far greater granularity than rule-based systems

Adaptive learning. Continuous optimization loops, where agents monitor outcomes, learn from variance and recalibrate execution parameters in real time

Interoperability. Cross-system coordination, where agents operate across platforms, APIs and data environments that previously required human intermediation

From decision support to decision execution

Earlier waves of enterprise AI such as predictive analytics, recommendation engines and natural language language-query interfaces supported human decisions. Agentic systems execute them. This transition is the defining feature of the AI-first GCC. It relocates the locus of execution from human effort to autonomous orchestration, with humans operating at a level focused on oversight and governance, and strategic direction.

language-query interfaces supported human decisions. Agentic systems execute them. This transition is the defining feature of the AI-first GCC. It relocates the locus of execution from human effort to autonomous orchestration, with humans operating at a level focused on oversight and governance, and strategic direction.

Research evidence suggests that approximately 24% of enterprise tasks are fully automatable with current technology, while another 42% can be significantly augmented, with meaningful productivity uplift from human-AI collaboration. Interpreted conservatively, this implies that nearly two-thirds of enterprise task time is structurally redesignable. The GCC, as the execution center of the enterprise, sits at the heart of this opportunity.

The agentic paradigm also has important implications for how GCCs think about their strategic contribution. A GCC that deploys Agentic AI is not simply doing existing work faster. It is changing the content of its activities, the structure of its linkages with the broader enterprise and the governance of its execution layer. These are the dimensions of business model innovation—and they are increasingly within reach of GCC leadership.



A man with a beard and glasses is wearing a VR headset. He is looking out a window at a city skyline. The scene is dimly lit, with light coming from the window and a ceiling light fixture.

CHAPTER 4

GCCs as enterprise AI launchpads

GCCs possess a unique structural configuration that positions them as natural AI acceleration hubs. This is not merely a matter of scale. It is a matter of the specific combination of conditions that makes GCCs, more than any other organizational unit, suited to lead enterprise AI transformation.

The relevant conditions span multiple dimensions:

Talent density: Concentrated cross-functional talent with deep process knowledge across finance, engineering, analytics, operations and compliance domains

Governance maturity: Mature governance frameworks developed over years of audit, risk and regulatory engagement, which provides the foundation for Responsible AI deployment

Data architecture: Enterprise-wide data access, often unique to GCCs, which enables cross-functional AI models that cannot be built from within siloed business units

Engineering scale: Scalable engineering capacity capable of industrializing prototypes into enterprise-grade deployments

Experimentation infrastructure: Cost-effective experimentation environments that allow hypothesis testing at lower risk than headquarters-led programs

The Center of Gravity principle

Fragmented AI deployment—where each function builds its own stack independently—creates redundancy, increases governance risk and dilutes institutional learning. A GCC operating as the enterprise's AI Centre of Gravity develops reusable agentic frameworks, enforces architectural discipline and accumulates a strategic asset that compounds over time. This is how a delivery utility becomes a global autonomous intelligence center.

The strategic shift implied here is significant. It requires GCC leadership to reconceive their mandate— not merely to deliver services efficiently but to own and evolve the enterprise's AI execution architecture. This means assuming responsibility not just for today's workflows but for the platforms, standards and capabilities on which tomorrow's operating model will run.

Leadership support for this mandate is not optional. It is the enabling condition. GCCs that achieve genuine AI-first transformation consistently exhibit strong headquarters sponsorship, clear innovation mandates, and explicit

governance frameworks that allow for experimentation while maintaining accountability. Budget follows ambition – but only when ambition is supported by the architecture.

Notably, the path to this position is not identical for all GCCs. Some will lead through aggressive technology adoption and GenAI capability building. Others will leverage operating model maturity and digital ownership to scale transformation more systematically. Still others will rely on innovation-driven culture and leadership sponsorship as the primary catalyst. These are not competing approaches – they are different configurational pathways to the same strategic destination.



CHAPTER 5

Industry archetypes and impact profiles

The AI-first transformation opportunity is not uniform across industries. Task structure, data maturity, regulatory complexity and workflow architecture shape the exposure profile and realization pathway. The following archetypes illustrate the range of impact under conservative assumptions.

Financial Services GCC

Financial Services GCCs benefit from high automation exposure in reconciliation, reporting and compliance monitoring. The structured, rule-dense nature of financial workflows creates an ideal environment for agentic deployment. Conservative productivity uplift in transactional areas can reach 30% to 35%, with additional value from compliance acceleration via explainable AI frameworks that reduce audit cycle times and regulatory reporting costs.

Manufacturing and Supply Chain GCC

Manufacturing GCCs exhibit strong working capital and logistics optimization potential. Agentic systems deployed across demand sensing, dynamic replenishment and API-connected logistics coordination can unlock measurable improvements in inventory productivity and significantly reduce expediting costs. The real-time nature of supply chain data creates a natural environment for continuous agent-driven recalibration.

Technology GCC

Technology GCCs are positioned to capture the highest engineering productivity uplift – 15% to 25% on developer output through AI-assisted software development, automated testing orchestration and code review acceleration. These centers also

have the highest potential for self-directed AI capability development, making them natural incubators for enterprise-wide agentic platforms.

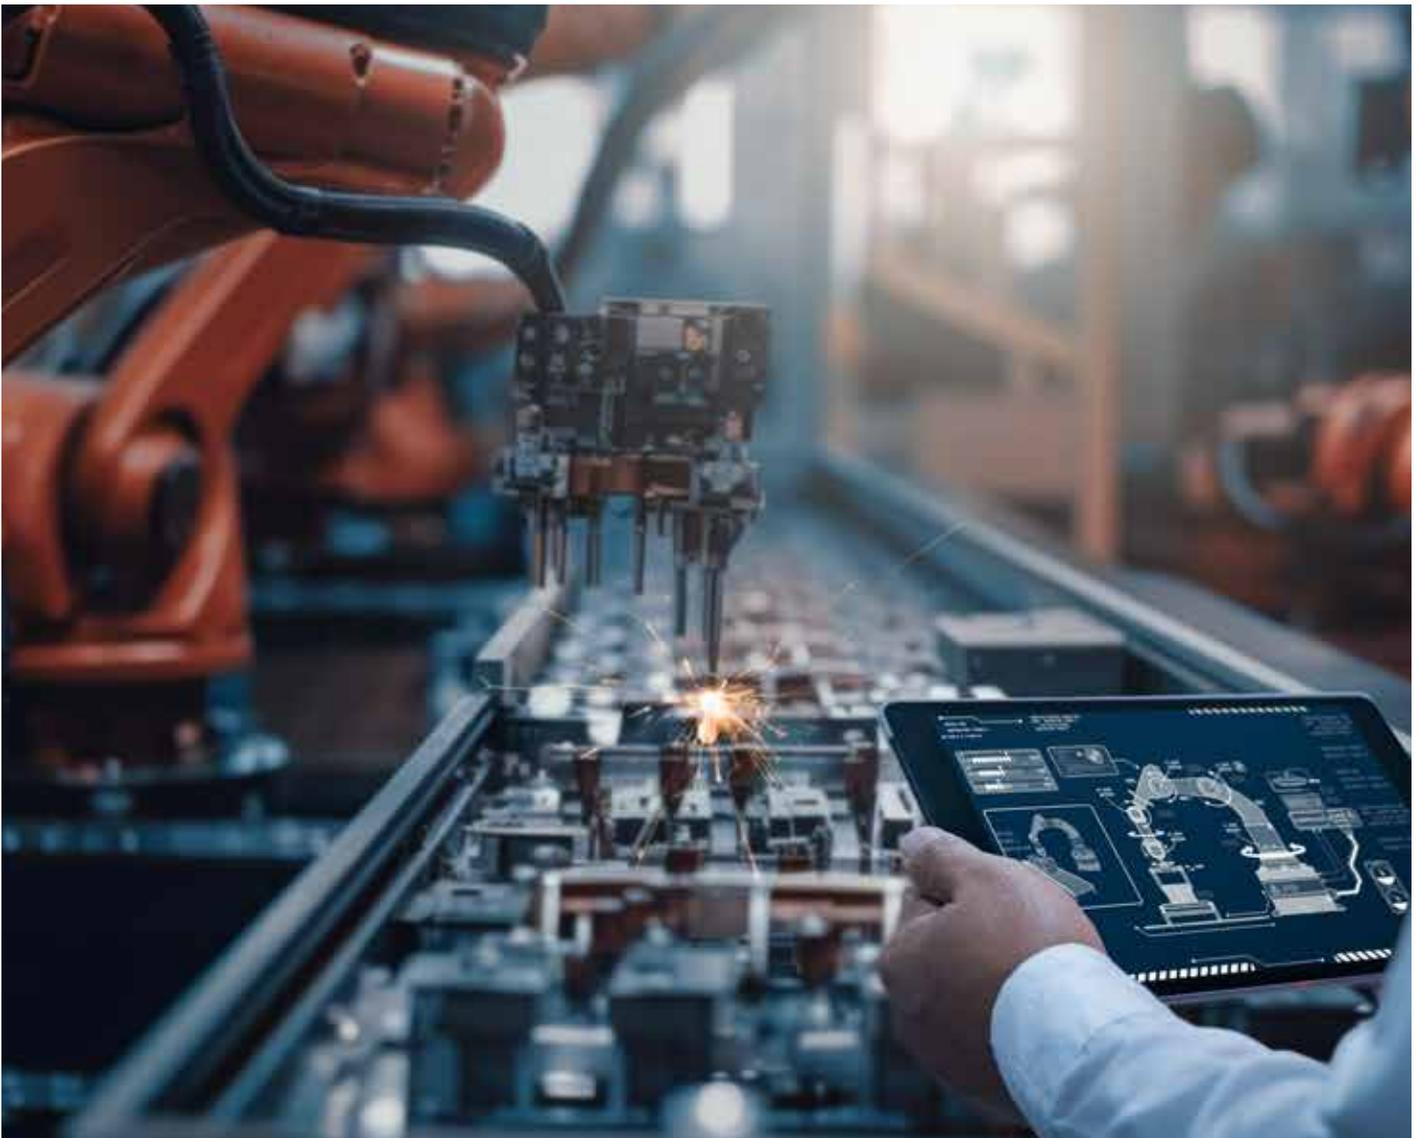
Healthcare and Life Sciences GCC

Healthcare and Life Sciences GCCs face a distinctive opportunity in regulatory documentation automation and clinical trial monitoring.

Documentation-intensive workflows, typically characterized by high variability

and specialist bottlenecks, can achieve 20% to 30% cycle-time reductions with well-governed agentic deployment. Data privacy and regulatory constraints require particularly rigorous governance architecture, but they do not preclude transformation.

Conservative blended productivity uplift across industry archetypes consistently exceeds 20% – before any revenue or innovation multiplier effects are included.





CHAPTER 6

Case illustration- AI-first Finance & Supply Chain GCC

The following illustration grounds the strategic argument in operational specifics. It depicts a representative transition from a traditional GCC model to an AI-first architecture in two adjacent domains: financial operations and supply chain management.

The starting point: Traditional model

A 250-FTE GCC supporting finance and supply chain operations. Workflows are characterized by manual reconciliations, weekly demand planning cycles, exception backlogs accumulating over multi-day periods and significant human intermediation at each process handoff. The center is well-run by conventional standards, but its operating model is structurally misaligned with the speed and complexity of the business environment it serves.

The finance close agent

In the AI-first model, the financial close process is governed by an autonomous agent that performs continuous ledger ingestion, real-time anomaly detection, reconciliation execution across accounts and draft financial statement preparation – without waiting for scheduled batch runs or human hand-offs. The finance close agent does not replace human judgment; it executes the structured components of the process at machine speed, surfacing only genuine exceptions that require human assessment.

The supply chain orchestration agent

The supply chain agent operates across real-time demand-signal ingestion, dynamic replenishment logic calibrated to lead times and safety stock parameters and API-based logistics optimization

across carrier networks. It replaces weekly planning cycles with continuous recalibration, reducing the latency between the demand signal and the fulfillment response from days to minutes.

Conservative outcomes under AI-first architecture

35% to 40% cycle-time reduction in financial close and supply chain planning processes

30%-35% 30% to 35% reduction in manual effort, with redeployment toward governance, analytics and strategic advisory roles

15% to 20% improvement in forecast accuracy, reflecting the advantage of continuous AI-driven sensing over periodic human analysis

What this illustration demonstrates is not merely an efficiency gain. It is a restructuring of the activity system: new activities have been introduced (continuous AI-driven monitoring and execution), existing linkages have been redesigned (from sequential hand-offs to event-driven orchestration), and the

governance of execution has shifted (from human-controlled batched processes to agent-controlled continuous processes with human oversight). Human roles do not disappear – they migrate toward governance, analytics and strategic advisory functions. This is the hallmark of genuine operating model reinvention.



CHAPTER 7

The four pillars of an AI-first GCC

Building an AI-first GCC is not a matter of deploying a platform. It requires simultaneous, interlocking investment across four structural dimensions. No single pillar is sufficient in isolation. The evidence from leading GCC transformations consistently points to the importance of alignment across all four – and the failure modes that occur when one is neglected.

Agentic-native talent and culture

Systematic reskilling toward AI supervision, workflow design, agent performance oversight and governance roles. Human effort becomes exception-based rather than repetitive, anchored in judgment, creativity and accountability

Agentic platforms and architecture

Hybrid model stacks, API-first orchestration layers and modular microservice architectures that enable interoperability, vendor independence and rapid capability extension across all enterprise domains

Governance and accountability frameworks

Explainability-by-design, continuous bias monitoring, human override controls, and audit trails that embed Responsible AI into every workflow layer—not as a compliance overlay but as a structural architecture

Operating model redesign

Event-driven orchestration replacing linear approval hierarchies. Processes are redesigned around AI intelligent logic: continuous ingestion, real-time decisioning and dynamic recalibration – shifting the GCC from a delivery node to an intelligence hub

The interdependence of these pillars is not incidental; it is the point. A center that builds exceptional agentic platforms without embedding governance is likely to face compliance risk and stakeholder loss of confidence. A center that invests in talent reskilling without redesigning its operating model may see AI-capable employees return to legacy workflows. A center that redesigns its operating model without committing to a platform architecture may build process blueprints it cannot execute.

This configurational logic – that outcomes depend on combinations of conditions rather than any single investment – has important practical implications. It means that the diagnostic question for GCC leadership is not simply 'what is our AI maturity?'; it is 'which of these pillars are we strongest in, which are most underdeveloped and what is the minimum bundle we need to activate to unlock sustainable transformation?'

The answer will differ by GCC. Centers with exceptional governance maturity may find that Pillar 3 accelerates the others. Centers with strong engineering depth may find that Pillar 2 provides the fastest path to demonstrated value. Centers with a powerful innovation-driven culture embedded by leadership may find that Pillar 1 becomes the multiplier that makes all other investment yield returns. There is no single recipe – but there are conditions that must be jointly present for transformation to scale beyond pilots.





CHAPTER 8

Governance, risk & Responsible AI architecture

AI-first transformation introduces a risk surface that differs qualitatively from conventional enterprise technology risk. The autonomous nature of agentic systems – their capacity to perceive, decide and act without constant human intermediation – creates new exposure dimensions that demand a deliberate architectural response.

The relevant risk categories include:

Reliability

Model reliability risk, including hallucination, reasoning errors and out-of-distribution performance in workflows where AI agents encounter contexts outside their training distribution

Security

Cybersecurity and adversarial risk, where AI systems operate with broad API access and tool-use capabilities present an expanded attack surface

Compliance

Data privacy and regulatory risk, particularly in jurisdictions with evolving AI governance frameworks and sectoral data protection mandates

Change Management

Organizational change risk, including resistance from a workforce population uncertain about role evolution and from senior stakeholders unfamiliar with agentic operating models

Responsible AI architecture addresses these risks not through restriction but through design. Key architectural principles include:

Principles of Responsible agentic architecture

- ▶ Human-in-the-loop supervision at decision-making points with material consequence
- ▶ Continuous monitoring dashboards with anomaly detection and threshold-based escalation
- ▶ Modular deployment that constrains agent scope to well-defined domains, reducing blast radius in failure scenarios
- ▶ Robust data governance that controls what information agents can access, retain and transmit
- ▶ Explainability frameworks that produce auditable decision trails for regulatory review
- ▶ Executive sponsorship that signals the organizational seriousness of governance investment

Responsible AI is not a constraint on transformation – it is an enabler. GCCs that build a governance architecture early can develop a durable trust foundation with their parent enterprise stakeholders, regulators and workforce populations. This

trust is a strategic asset: it enables faster approval cycles for new agent deployments, supports talent retention in a landscape of legitimate concern about AI's role and positions the GCC as the enterprise's Responsible AI exemplar rather than its liability.



CHAPTER 9

Workforce transition and capability strategy

The AI-first GCC requires workforce redesign – not workforce reduction. This distinction is both analytically important and organizationally essential. The evidence from mature AI deployments consistently shows that the most significant constraint on value capture is not the capability of the AI systems. It is the capability of the human organization to work alongside them productively.

New roles emerge at the intersection of domain expertise and AI oversight:

- AI Workflow Architects who design, document and maintain the agent orchestration logic for complex business processes

- Prompt Engineers and Agent Designers who develop, test and iterate on the instruction sets and tooling configurations that govern agent behavior

- Agent Performance Supervisors who monitor live deployments, assess output quality and manage exception queues

- Responsible AI and Governance Leads who own the ethical, regulatory and audit frameworks across all agent deployments

Conservative transition modeling assumes that 60% to 70% of productivity gains are redeployed toward higher-value work rather than eliminated from the workforce. This is not wishful thinking – it is a strategic choice and one with strong business logic. GCCs that preserve institutional knowledge while upskilling toward AI intelligent roles build

transformation capacity that is difficult to replicate from scratch.

The alternative – aggressive workforce reduction – typically destroys the organizational memory and governance capability that successful AI deployment requires.

The GCC that invests in its people's capacity to govern, design and improve AI systems is likely to outperform the one that invests only in the systems themselves.

Innovation culture and leadership support are underrated multipliers in workforce transition. Centers that create psychological safety for experimentation, establish credible narratives about career evolution and provide structured learning pathways often consistently outperform

those that rely on top-down mandates without cultural investment. Culture is not a soft infrastructure – it is a strategic resource that differentiates centers capable of scaling AI transformation from those that accumulate pilots that never reach production.



CHAPTER 10

Implementation roadmap and phased financial model

Transformation at this scale does not happen in a single motion. The following roadmap reflects the sequencing logic that characterizes successful AI-first GCC implementations: beginning with rigorous diagnostic work, building architecture before scaling deployment and embedding governance as an ongoing discipline rather than a final gate.

The 12-month-plus transformation horizon is not a constraint – it is a discipline. Attempts to compress this timeline by skipping diagnostic or governance phases consistently produce fragile deployments that require costly remediation. The phased model also enables iterative financial validation: each phase should yield measurable value that funds the next, reducing dependence on a large capital commitment and building organizational confidence alongside technical capability.

Financial sensitivity modeling across the transformation horizon should test three scenarios: a base case (25% productivity uplift, gradual workforce redeployment), a conservative case (20% uplift, slower benefit realization) and an optimistic case (30% uplift, accelerated governance maturity). In all three scenarios, the net present value (NPV) case for action remains positive. The conservative case yields lower absolute returns but preserves the fundamental economic logic of transformation.

On the importance of sequencing

The most common failure mode in GCC AI transformation is not technological – it is organizational. Centers that deploy AI capabilities before establishing governance, workforce readiness, or operating-model clarity find themselves managing failures rather than scaling successes. The roadmap above is structured to enable each phase that creates the conditions in which the next can succeed.

| Phase | Focus Area | Horizon | Key Activities |
|---------|-----------------------------------|-------------------|---|
| Phase 1 | Diagnostic and exposure mapping | Months 1- 2 | Process decomposition; AI Automation-exposure scoring; workforce capability audit; technology landscape review |
| Phase 2 | Process re-architecture | Months 1- 2 | Workflow redesign around agentic orchestration; identification of high-value deployment domains; operating model blueprint |
| Phase 3 | Platform deployment | Months 2- 10 | Agentic-platform build and integration; pilot deployment in priority domains; human-AI workflow testing and calibration |
| Phase 4 | Governance embedding | Months 10- 12 | Responsible AI architecture in production; compliance integration; Workforce-transition support; change-management activation |
| Phase 5 | Scale and continuous optimization | Months 12 onwards | Enterprise-wide rollout; autonomous optimization loops; performance benchmarking; center of excellence maturation |

CHAPTER 11

Strategic implications for enterprise leadership

AI-first GCC transformation is a board-level strategic decision. It is not a technology program that can be delegated to a CIO or an automation initiative that can be owned by the Shared Services Director. Its implications span competitive positioning, cost structure defensibility, innovation velocity, regulatory resilience and workforce evolution.

For enterprise leadership, the strategic questions are clear:

Architecture

Is the GCC positioned as the enterprise's AI center of gravity, or is AI capability being built in fragmented, uncoordinated ways across business units?

Mandate

Does the GCC have the mandate, governance framework and leadership sponsorship needed to own the enterprise's agentic execution architecture?

Investment

Is transformation investment calibrated to the scale of the opportunity – or are resource constraints limiting the GCC to pilots that cannot achieve the scale needed for compounding returns?

People

Is the workforce transition being managed proactively, with structured capability pathways and cultural investment, or reactively, after resistance has already accumulated?

Governance

Is Responsible AI architecture being built into the design of every agentic deployment – or added as a compliance layer after the fact?

The competitive implications of inaction are asymmetric. Enterprises that redesign their GCC architecture now are building structural advantages – lower cost bases, faster decision cycles and deeper enterprise data integration – that will be difficult to close once established. The compounding nature of AI-driven productivity means that the gap between leaders and followers will widen over time, not narrow.

Enterprises that can define the next decade of competitive advantage are not those with the best AI tools. They are those with the best AI intelligent organizational architecture.

The GCC is not a supporting character in this story. It is the stage on which enterprise-level AI transformation will – or will not – be achieved. The leadership question is whether that stage is being built deliberately, with the architecture, capabilities and governance that the opportunity demands.



CHAPTER 12

Conclusion: Architecting the autonomous intelligence center

The AI-first GCC represents structural enterprise reinvention. It transforms execution layers, economic models and governance frameworks simultaneously. It demands new thinking about what a GCC is for, how its value is created and measured, and what combinations of organizational conditions enable it to contribute at the level of enterprise business model innovation – not merely operational efficiency.

Several conclusions emerge from the analysis presented in this paper.

First, the strategic opportunity is real and the economic case is robust. Under conservative modeling assumptions – with deliberate haircuts on realization rates and no credit for revenue or innovation upside – the financial returns on AI-first GCC transformation can be compelling. The NPV case is positive across scenario ranges. The cost of inaction increases over time.

Second, the pathway is not the same for every GCC. Different centers will arrive at AI-first maturity through different combinations of technology posture, digital ownership, innovation-driven culture, GenAI readiness, operating model maturity and transformation investment. Leadership should resist the temptation to benchmark against celebrated exemplars and instead diagnose their own organizational configuration – identifying which conditions are already present as strengths, which represent critical gaps, and what minimum bundle of aligned conditions must be activated to move from pilots to scaled transformation.

Third, governance is not a constraint – it is a multiplier. Organizations that embed Responsible AI architecture early build

trust, accountability and institutional legitimacy that enable faster, more ambitious deployments over time. Governance is not the ceiling of what AI-first GCCs can achieve. It is the foundation on which sustainable scale is built.

Fourth, people remain a critical variable. Technology platforms are necessary but

not sufficient. GCCs that can lead enterprise AI transformation are those that invest equally in their human capability infrastructure – reskilling toward oversight, design and governance roles; building innovation-driven cultures that embrace experimentation; and earning the organizational trust needed to lead, rather than merely host, enterprise-level change.

The strategic question is no longer whether to adopt AI. It is whether an enterprise is architecting autonomy deliberately, responsibly and at scale.

The organizations that embed Agentic AI into the foundation of their GCC strategy would be well placed to build scalable, sovereign, and resilient intelligence capabilities while preserving the governance discipline that enterprise stakeholders require. Under conservative assumptions, the case for action is already compelling. Under realistic assumptions, the opportunity is transformational.

The autonomous intelligence center is not a vision statement. It is an architectural choice. The question facing GCC leadership today is whether to make that choice deliberately – or to have it made by the competitive environment.

About this paper

This white paper was prepared for presentation to GCC executive leadership. The economic modeling applies conservative assumptions and intentionally excludes revenue and innovation-upside projections. All scenario estimates are illustrative and should be validated against organization-specific parameters. The strategic frameworks presented reflect current evidence on enterprise AI adoption, GCC operating model evolution and the organizational conditions associated with successful large-scale transformation programs.

Authoring team

Working group



Arindam Sen

Partner and GCC Sector Lead -
Technology, Media &
Entertainment and
Telecommunications, EY India



Sunil Venkatesh

Partner, Technology Consulting,
and Hyderabad GCC Sector Lead,
EY India



Mayank Gaur

Director
Business Consulting, EY India

EY GCC core team



Arindam Sen
Partner and GCC Sector Lead -
Technology, Media &
Entertainment and
Telecommunications, EY India



Ajay S Kamat
Partner,
Technology Consulting,
EY India



Ajay Sirikonda
Partner - Financial Services
Risk Consulting,
EY India



Alpana Dutta
Partner,
People Consulting,
EY India



Anand Mihir
Partner Leader - Financial
Services Risk Consulting,
EY India



**Balasubramanian
Manikandan**
Partner, Business Consulting
and Chennai GCC Leader,
EY India



Kunal Ghatak
Partner,
Business Consulting,
EY India



Manoj Marwah
Partner,
FS GCC Leader,
EY India



Radhika Saigal
Partner, FS Consulting
Leader, EY India



Rashmi P Kumat
Partner, Technology Consulting,
and West GCC Sector Lead,
EY India



**Ritika Loganey
Gupta**
Tax Partner,
EY India



Shiv Narain
Partner, GCC Consulting and
North & East GCC Sector Lead,
EY India



Subir Mehra
Partner and GCC Sector Lead -
Financial Services,
EY India



Sunil Venkatesh
Partner, Technology Consulting,
and Hyderabad GCC Sector
Lead, EY India



Tiffy Isaac
Partner,
Risk Consulting,
EY India



Vijay S Bhaskaran
Partner,
Business Consulting,
EY India



About AMCHAM

The American Chamber of Commerce in India

AMCHAM India is the leading apex chamber of U.S. industry in India. Established in 1992, AMCHAM has over 400 U.S. companies as members and plays a pivotal role in fostering strong ties between the U.S. and India. The incumbent U.S. Ambassador to India is the Honorary President of AMCHAM. The chamber enjoys a close relationship with the U.S. Embassy and complete support in fulfilling its objectives. Country Heads of leading U.S. companies constitute the elected national executive board. The chamber's mission is to assist member companies to succeed in India through advocacy, information, networking and business support services. Headquartered in New Delhi, AMCHAM extends its influence through regional chapters in Bengaluru, Chennai, Hyderabad, Kolkata, Mumbai and Pune.



Our offices

Ahmedabad

22nd Floor, B Wing, Privilon
Ambli BRT Road, Behind Iskcon Temple
Off SG Highway
Ahmedabad - 380 059
Tel: + 91 79 6608 3800

8th Floor, Building No. 14A
Block 14, Zone 1
Brigade International Financial Centre
GIFT City SEZ
Gandhinagar - 382355, Gujarat
Tel +91 79 6608 3800

Bengaluru

12th & 13th Floor
"UB City", Canberra Block
No.24 Vittal Mallya Road,
Bengaluru - 560 001
Tel: + 91 80 6727 5000

Ground & 1st Floor
11, 'A' wing
Divyasree Chambers
Langford Town,
Bengaluru - 560 025
Tel: + 91 80 6727 5000

3rd & 4th Floor
MARKSQUARE
#61, St. Mark's Road
Shantala Nagar,
Bengaluru - 560 001
Tel: + 91 80 6727 5000

1st & 8th Floor, Tower A
Prestige Shantiniketan
Mahadevapura Post
Whitefield,
Bengaluru - 560 048
Tel: + 91 80 6727 5000

Bhubaneswar

8th Floor, O-Hub, Tower A
Chandaka SEZ, Bhubaneswar,
Odisha - 751024
Tel: + 91 674 274 4490

Chandigarh

Elante offices, Unit No. B-613 & 614
6th Floor, Plot No- 178-178A
Industrial & Business Park, Phase-I
Chandigarh - 160 002
Tel: + 91 172 6717800

6th & 7th Floor, A Block,
Tidel Park, No.4, Rajiv Gandhi Salai
Taramani, Chennai - 600 113
Tel: + 91 44 6654 8100

Delhi NCR

Aikyam
Ground Floor
67, Institutional Area
Sector 44w, Gurugram - 122 003
Haryana
Tel: + 91 124 443 4000

3rd & 6th Floor, Worldmark-1
IGI Airport Hospitality District
Aerocity, New Delhi - 110 037
Tel: + 91 11 4731 8000

Hyderabad

THE SKYVIEW 10
18th Floor, "SOUTH LOBBY"
Survey No 83/1, Raidurgam
Hyderabad - 500 032
Tel: + 91 40 6736 2000

Jaipur

9th floor, Jewel of India
Horizon Tower, JLN Marg
Opp Jaipur Stock Exchange
Jaipur, Rajasthan - 302018

Kochi

9th Floor, ABAD Nucleus
NH-49, Maradu PO
Kochi - 682 304
Tel: + 91 484 433 4000

Kolkata

22 Camac Street
3rd Floor, Block 'C'
Kolkata - 700 016
Tel: +91 33 6615 3400

6th floor, Sector V,
Building Omega, Bengal Intelligent
Park, Salt Lake Electronics Complex,
Bidhan Nagar, Kolkata - 700 091
Tel: +91 33 6615 3400

14th Floor, The Ruby
29 Senapati Bapat Marg
Dadar (W), Mumbai - 400 028
Tel: + 91 22 6192 0000

5th Floor, Block B-2
Nirlon Knowledge Park
Off. Western Express
Highway, Goregaon (E),
Mumbai - 400 063
Tel: + 91 22 6192 0000

3rd Floor, Unit No.301
Building No.1, Mindspace-Gigaplex
IT Park, MIDC, Plot No. IT-5
Airoli Knowledge Park
Airoli West, Navi Mumbai - 400 708
Tel: + 91 22 6192 0003

Altimus, 18th Floor
Pandurang Budhkar Marg
Worli, Mumbai - 400 018
Tel: +91 22 6192 0503

Pune

C-401, 4th Floor
Panchshil Tech Park, Yerwada
(Near Don Bosco School)
Pune - 411 006
Tel: + 91 20 4912 6000

10th Floor, Smartworks
M-Agile, Pan Card Club Road
Baner, Pune - 411 045
Tel: + 91 20 4912 6800

Ernst & Young LLP

EY | Building a better working world

EY is building a better working world by creating new value for clients, people, society and the planet, while building trust in capital markets.

Enabled by data, AI and advanced technology, EY teams help clients shape the future with confidence and develop answers for the most pressing issues of today and tomorrow.

EY teams work across a full spectrum of services in assurance, consulting, tax, strategy and transactions. Fueled by sector insights, a globally connected, multi-disciplinary network and diverse ecosystem partners, EY teams can provide services in more than 150 countries and territories.

All in to shape the future with confidence.

EY refers to the global organization, and may refer to one or more, of the member firms of Ernst & Young Global Limited, each of which is a separate legal entity. Ernst & Young Global Limited, a UK company limited by guarantee, does not provide services to clients. Information about how EY collects and uses personal data and a description of the rights individuals have under data protection legislation are available via ey.com/privacy. EYG member firms do not practice law where prohibited by local laws. For more information about our organization, please visit ey.com.

Ernst & Young LLP is one of the Indian client serving member firms of EYGM Limited. For more information about our organization, please visit www.ey.com/en_in.

Ernst & Young LLP is a Limited Liability Partnership, registered under the Limited Liability Partnership Act, 2008 in India, having its registered office at Ground Floor, Plot No. 67, Institutional Area, Sector - 44, Gurugram - 122 003, Haryana, India.

© 2026 Ernst & Young LLP. Published in India.
All Rights Reserved.

EYIN2603-020
ED None

This publication contains information in summary form and is therefore intended for general guidance only. It is not intended to be a substitute for detailed research or the exercise of professional judgment. Neither EYGM Limited nor any other member of the global Ernst & Young organization can accept any responsibility for loss occasioned to any person acting or refraining from action as a result of any material in this publication. On any specific matter, reference should be made to the appropriate advisor.

RS

ey.com/en_in

