

Now that (almost) Everyone has AI

The question is no longer whether your company uses AI. It is whether your business model gets stronger or weaker every time AI improves, and how IT service providers can best respond to this.

September 2026

By Rajiv Gupta, Akash Bhatia, Mandeep Kohli, Karan Chadha, Jade Leet, and Cara Castellana Kreisman



1. The Widening Wedge

Most companies treat AI improvement as a rising tide that lifts everyone. It is not. It may be better understood as wedge that widens roughly every six weeks with every new model release.

On one hand: There are service providers who accumulate operational knowledge – process depth, exception patterns, domain intelligence, and client trust, giving themselves a chance to become more valuable with every model upgrade. As AI improves, they become faster, cheaper, and more deeply embedded. Their advantage potentially compounds.

On the other: There are companies whose current offerings get commoditized by the same improvements. Often, a new model release erodes what they sold last quarter. And many capabilities that were differentiated just six months ago becomes available as a commodity API. The gap seldom closes. It almost always widens.

The cost of building custom software is falling sharply. What took a team of ten developers six months can now be prototyped in days. This suggests a structural reset of the input cost of delivery rather than an incremental improvement.

Per-seat pricing is breaking down because agents do not log in. The shift introduces its own cost structure: token consumption at scale can be substantial, but the unit of pricing is changing from headcount to work done. A major frontier model is now launched roughly every six weeks. Between 2023 and mid-2026, the four leading labs (OpenAI, Anthropic, Google DeepMind, and Meta) collectively released more than 30 major frontier model versions. The cost of GPT-4-class inference has fallen approximately 60x in under two years. Open-weight models are also catching up quickly, and because they can be run in-house cheaply, they push the cost of capable AI down further still. Capabilities that were cutting-edge eighteen months ago are available as commodity APIs. However, this does not mean all models are equal. Frontier capabilities remain meaningfully differentiated, and enterprises pay a premium for them.

For technology services companies, this looks unlike any earlier wave.

Cloud, digital transformation, and analytics were all largely additive. They gave infrastructure teams new platforms, integrators larger backlogs, and outsourcers more processes to manage. Each wave expanded the addressable market without undermining the model underneath. However, AI, especially agentic AI, is different. It threatens to replace the work itself while creating new work previously not even imaginable.

Until recently, volume expansion offset price compression. That buffer is running out; and when it goes, the terminal value of the current model becomes the question.

The leadership layer is underestimating the risk: In a recent survey of 467 senior executives at \$1B+ enterprises, 87% said their CEO and board have significant gaps in understanding AI investment risk. And 83% said their leadership does not adequately understand that the risk of underinvesting may be existential.

A small number of companies – mostly AI-native players and outcome-priced software firms – are already on the right side of the wedge. Their models are designed so that every AI improvement makes them faster, cheaper, and more deeply embedded.

For everyone else, the penalty is real, and it compounds with every quarter of delay.

2. The Incremental Playbook is Running Out of Room

Most players have run the rational playbook: rebrand the data practice as an AI offer, sign a hyperscaler deal, pilot a few agents, and put “agentic” in the investor deck. These moves signaled relevance, but they are no longer sufficient.

The clients who once rewarded AI ambition are now increasingly measuring execution, and the client who praised your AI roadmap in the QBR may still ask for a rate reduction in the next contract cycle. The gap between what providers claim and what enterprises experience is becoming impossible to ignore.

The buyer’s ask is shifting.

Across 70+ CXO interviews with enterprise leaders across industries, there were six recurring asks: (1) clear business value tied to P&L lines, (2) context specificity rooted in how the work actually runs, (3) integration into the existing stack, (4) data sovereignty and governance, (5) in-house control and flexibility, and (6) proof of value with outcome ownership. The common underlying theme: who owns the result.

A sea of sameness has made differentiation difficult

We analyzed 40+ providers across five archetypes: (1) Hyperscalers, (2) SaaS platforms, (3) BPO firms, (4) Systems integrators, and (5) AI-native boutiques, and pulled more than 100 positioning claims from their websites, investor presentations, and earnings calls. We wanted to know one thing: is there anything that actually sets these players apart?

The answer, mostly, is no.

The exhibit 1 indicates this visually. The players are lined up on the top row. Down the side are the themes these providers talk about: multi-agent orchestration, domain-specific agents, autonomous task execution, among other features. The same feature appears across the market as “multi-agent collaboration”, “connected agent workflows”, and “orchestrated agent ecosystems” – different words, essentially the same concept. One provider frames it as a platform play, another as an architecture, a third as a delivery model – but read past the label and they are describing much the same thing.

The same holds for “autonomous task execution,” the idea of agents completing work end-to-end. One provider calls it an “autonomous workforce”, another an “AI-powered execution layer”, a third simply “agents that act and deliver end-to-end”. The framing shifts; the claim does not. Across much of the exhibit the pattern repeats: the same themes, in slightly different words, claimed by nearly everyone. That is what a crowded market looks like. The capabilities that would actually set a provider apart are barely spoken of. In a market this crowded, differentiation comes from authenticity, knowing exactly where you play and being able to demonstrate it credibly.

EXHIBIT 1

Sea of sameness: Where differentiation exists vs. where it does not

Keyword	Multi-Agent Orchestration	Domain-Specific Agents	Autonomous Task Execution	Agent Arch & Control Planes	Real-Time Intelligent Ops	Full-Stack AI Platform
Company 1	Open A2A protocol		Computer use capabilities to interact	Vertex AI Agent Builder	Real-time data to adapt	Full stack approach to AI innovation
Company 2	Bedrock multi-agent collaboration	Security Agent; DevOps Agent	Automating tasks	Bedrock AgentCore	Real-time agent responses	
Company 3	Multi-agent frameworks orchestration	M365 Analyst; M365 Researcher	Fully autonomous actions; asynchronous coding	Foundry Azure AI		
Company 4			Operator browses web autonomously; task completion	Responses API; agent-native APIs		
Company 5			Agency [to] act on its own	Model Context Protocol (MCP)		
Company 6	A2A protocol; Joule multi-agent workflow	Joule Agents across customer experience	Joule agents... decide the action	Joule Studio; AI Agent Hub		Broadest suite of applications in industry
Company 7	A2A comms.; Agentforce 360	300+ pre-built industry specific agents; Agentforce	Limitless digital workforce of AI	Agentforce Platform; Agentforce Studio	Real-time data; agents that adapt	Agentforce 360; connective tissue
Company 8	AI Control Tower; agent orchestrator	AI agents across ITSM, ITOM	Autonomous, self-driven systems; entire-touch automation	AI Control Tower; Now Platform	Real-time workflow intelligence	One platform, one architecture
Company 9	GenWizard, SynOps and AI Refinery	Prebuilt industry agents and workflows	AI that continuously learns, acts	AI Refinery		
Company 10	Multi-agent orchestration	AI agents tailored to specific needs	Domain-aware agents to work alongside	AI Fabric		

Continued ...

Keyword	Multi-Agent Orchestration	Domain- Specific Agents	Autonomous Task Execution	Agent Arch & Control Planes	Real-Time Intelligent Ops	Full-Stack AI Platform
Company 11	Multi-agent systems	Next Best Actions (NBA)	Agentic AI to autonomously resolve			
Company 12	Multi-agent Orchestrate	Pre-built agents covering HR, sales, etc.	AI agents that act autonomously	Semantic Control Plane		
Company 13	AI Force 2.0; full-stack AI ecosystem	Agentic enterprise solutions across HC	Zero-touch automation and agentic solutions	AI Force 2.0; AI Foundry		Full-stack AI ecosystem
Company 14	Ai360; multi-agent workflows	More than 200 AI agents	Agentic and assistive AI			
Company 15	Agent Foundry; multi-agent enterprise scale		Agentic and autonomous enterprise AI		Real-time decisioning	
Company 16	250+ pre-built agents	ClaimsAssist.ai; FCRM; healthcare	Agent Studio; no-code autonomous agent depl	Governance Hub; 40+ specialized models		Agent Studio + Governance Hub
Company 17	SASVA 4.0 Brain; visual orchestration	BFSI-focused; SASVA for banks	40-70% efficiency in back/middle office	SASVA Brain; deterministic guardrails		SASVA 4.0
Company 18	Forge-X; AgentSphere	P&C; insurance; travel	EvolveOps.AI autonomous IT ops	Forge-X; SDLC + PDLC agentic platform	EvolveOps.AI / MTTD/MTTR cut 60%	Forge-X + Quasar GenAI Central + AgentSphere
Company 19	C3 Type System; agentic orch layer	Reliability; product; SCM; energy mgmt	Agentic Process Automation; NL to workflow	C3 AI Studio; Model-driven arch.	IoT: real-time sensor analytics	End-to-end develop / operate
Company 20	AIP Ontology; 12-layer agentic arch	Defense; supply chain; underwriting	AIP Automate	Ontology-centric; 4 memory types	Foundry + AIP	Apollo; full OS for enterprise AI


We ran three deep analyses on the underlying claims. We undertook semantic clustering, grouping the terms that really mean the same thing, to see how many distinct ideas there actually are. We then measured signal velocity to understand how loudly each idea is being talked about. Finally, we used market echo saturation to do a count of the number of players who claim each idea.

1. Semantic Clustering: We mapped keywords that are conceptually co-dependent. We clustered the terms that must travel together into a single, non-overlapping set of themes. The 100+ keywords in circulation collapsed into just 20+ distinct plays, meaning much of the market's apparent diversity of language is mostly illusion.

What looks like a hundred different pitches is really about twenty ideas, repeated in different words. 'Autonomous task execution,' 'agentic operations,' and 'self-learning agents' all describe the same idea. 'Multi-agent orchestration,' 'agent architectures,' and 'protocol standardization' all describe the same thing: how agents connect and coordinate. That leaves buyers with a hard truth: the words can't tell providers apart anymore. The clustering also revealed something the vocabulary obscures: beneath the noise, only a handful of plays remain genuinely unclaimed: end-to-end accountability, outcome-based pricing, and agentic managed services. These are the spaces where a credible position can still be built. So next, we looked at how loudly each idea is being pushed, and by how many. See Exhibit 2

EXHIBIT 2

Semantic clustering: 100+ keywords converge into 23 distinct plays




2. Signal Velocity: We measured how much each of the 20+ plays was being talked about – across earnings calls, press releases, and BCG Executive Tracker data – with the conversations radiating outward from a center of intensity. A small group of themes populated that center, dominating many earnings calls. “Multi-agent orchestration” was the loudest of them all. A hyperscaler, a CRM platform,

and two services firms all led with it, using almost identical language. Closely behind were “governance and observability” and “measurable outcomes”, the terms that almost every provider uses to signal reassurance. Further out of the center were the ideas losing steam: legacy modernization, low-code, and autonomous enterprise. See Exhibit 3

EXHIBIT 3

Signal Velocity: Loudest conversations at the center, fading echoes at the edge



What the rings mean

We scored each of the 23 plays on the velocity of recent market discourse - across earnings calls, press releases, BCG Executive Tracker data, etc. Conversations radiate outward from a centre of intensity.

Highest - Accelerating sharply At the core: where discourse is fastest; these conversations dominate every earnings call	High - Gaining momentum Close to the centre: rising fast, attention compounding quarter after quarter	Medium - Steady discourse Mid-distance: present in the conversation but not gaining velocity	Fading - Losing steam Outer edge: discourse retreating as agentic categories supersede them
---	--	---	--

3. Market Echo Saturation: Finally, we measured how crowded each theme already was. We measured how many top-tier players were claiming each theme, and how loudly did they do so. The idea was to figure out where the room is jammed and where whitespace still exists.

The most crowded themes are the ones nearly everyone claims, so loudly that no single voice stands out. These themes are multi-agent orchestration, governance, adaptive agent intelligence, and responsible AI. Take the example of responsible AI: one provider calls it a “responsible

AI framework”, another “trustworthy and explainable AI”, a third “built-in guardrails”, a fourth “AI governance by design”, and a fifth folds it into “governance and observability” inside their agent platform. The same promise – trust, safety, and control – communicated in five different ways, with almost none explaining what they actually do differently. The themes that would actually set a provider apart sit almost empty: end-to-end accountability and agentic managed services are claimed by only a few. The whitespace is real, and hardly anyone is occupying it. See Exhibit 4

EXHIBIT 4

Market Echo Saturation: Where the room is jammed vs. where whitespace exists

Market Echo Saturation - What are players talking about constantly

Measuring how loudly and how many top-tier players are already claiming each theme - identifying where the room is jammed vs where whitespace exists



The analyses reveal five insights:

- **Most players are crowded onto the tech story:**

Almost all players are competing on the same technical ground: platforms, orchestration, autonomy. Over 90% of them talk about multi-agent orchestration, autonomous task execution, and agent architectures, which makes it nearly impossible to tell them apart.

- **Players avoid what buyers most ask for:**

Messaging leans on technical strengths and product capability, while the things enterprises actually asked for – proof it fits how the work runs, control over their own environment, and someone accountable for the result – go largely unspoken.

- **Few address how the work itself changes:**

Only a few players explain what the introduction of AI ends up doing to the workforce, how people and AI can work together, or what kind of reskilling may be needed. Fewer than 20% address cultural and operating-model change head-on, even though 45%+ of enterprises say this is primarily what is holding them back.

- **Data talk is everywhere but says nothing:**

Everyone says data matters, but almost no one says anything distinctive about it. More than 70% stress its importance, yet fewer than 25% talk about differentiated semantic graphs, infrastructure, or data layers.

- **Sameness is most visible where the real work sits:**

Service providers cluster hardest around the familiar areas: workflow transformation, domain-led app development, responsible AI frameworks. About 85% of their messaging sits there. Very few take on the harder, outcome-oriented questions like responsible scaling in regulated settings, human-in-the-loop operating models, and getting paid for results.

What separates providers now is not what they say, but what they can prove. When everyone claims the same ground in the same words, the words stop meaning anything.

The leap needed is not just in marketing language but also in positioning authentically and building the operating model to prove it.

The market is loud on what AI can do,
and quiet on who owns the outcome.

Closing that gap is not a messaging problem.
It is a business model problem.

3. Where the Shift is Real, and Why the Outcome Model Compounds

The widening wedge does not play out uniformly across tech services. It hits different segments at different speeds.

The shift is most visible in outsourced operations:

BPO, managed services, and customer experience.

This is where new business models are emerging fastest.

It looks different in software engineering and application development, the area drawing the most attention.

Here, there has been an aggressive disruption to the old headcount-based model: AI is compressing billable hours faster than almost anywhere else. But turning that speed into clean, owned outcomes is harder, because code quality is difficult to measure and easy to inflate. So, while the pressure on pricing is intense, the move to genuine outcome ownership is less settled than in operations, where results are easier to define and stand behind. The reason operations is ahead is not accidental; it follows directly from the economics of how enterprises spend on technology.

Operations sit at the intersection of repeatable processes, measurable outcomes, and large labor spend. When an AI agent can handle a customer inquiry, process an invoice, or resolve an IT ticket, the economic logic of the traditional model – large headcount, linear cost structure, and FTE-based pricing – comes under direct pressure.

EXHIBIT 5

Access/Effort vs. operator model, the structural divergence

A new category is coming up to fill the gap and claim the larger prize

For every \$1 spent on software, enterprises spend \$6 running it; that \$6 – is where durable advantage will come from



The exhibit 5 above captures the structural divide in a single ratio. For every \$1 an enterprise spends on software, it spends roughly \$6 running it – and that \$6 is where durable advantage now forms.

There are only two ways to capture it, and they age in opposite directions. The access and effort model charges for inputs – seats, licenses, time and materials – regardless of whether the outcome improves or not. Every AI upgrade threatens the value of what was just sold, because the buyer can now get more from the tool directly. The operator model charges for the work done – which means revenues are tied to the results delivered. Every upgrade makes the operator faster and cheaper on the same cost base.

One model fights technology. The other compounds with it. And the prize is far larger than it looks: the services and operations market sitting on top of that \$6 is several times the size of the software market beneath it.

Why this matters boils down to the question of what commoditizes and what doesn't.

As AI advances, execution gets cheaper. Anyone can reach the same models, the same agents, and the same raw capability. What does not commoditize is the ability to own a result credibly in a specific domain, with clients who trust you to deliver it: knowing which process to automate, managing the

exceptions the model was never trained on, holding the line under an SLA when something breaks, and getting smarter about the process with every cycle. Execution becomes cheap. Owning and compounding the outcome becomes the scarce, premium thing.

Three things separate the operators who compound from the providers who merely promise:

1. They get stronger when AI improves

Every model upgrade lowers cost-per-transaction and deepens the process map. Access – and effort-based models face the opposite result – an upgrade eats into what they sold last quarter.

2. They stay permanently, not project-to-project

Models drift, pipelines break, regulations change; but the operator who stays owns the compounding advantage. Everyone else resets to virtually zero.

3. They own the work, don't just enable it

Anyone can write an outcome into a contract. The edge lies in being operationally built to deliver it – living in the client's stack, working in the trenches, and keeping the operations running even as the models underneath keep changing. This is what sets the new AI-native services players apart: they don't hand over capability, they run the work and answer for the result.

The companies being built from scratch around this logic are already delivering tangible results. This new class of companies has no FTE pyramid to feed, no per-seat contracts to defend, no legacy delivery models to protect, and no revenue guidance built around the old model.

One AI-native customer experience platform charges only when its AI resolves an interaction – not per seat, not per head. It reached more than \$150M in recurring revenue within eight quarters on that model. Rather than handing over a tool, it embeds its forward-deployed engineers (FDEs) inside the operation, building and adapting the AI against the real workflow until the outcome shows up. For these players, the economics run the other way: when the AI improves, they earn more on the same cost base, not less.

Consider what the shift looks like in practice: a major US healthcare payer was running prior-authorization through a traditional labor-intensive model. This involved hundreds of analysts, manual review workflows, high error rates, and cycle times measured in days. An AI operations startup, founded by former Palantir engineers and now valued at close to \$2B, embedded its platform directly into the operation. The AI did not replace the analysts. It restructured the workflow around them, automating routine determinations, surfacing the right clinical context for complex cases, and learning from every exception. Within twelve months, the operation delivered \$23M in annual savings, not from headcount reduction alone, but from fundamentally changing how the work gets done. The commercial model was tied to measurable outcomes, not to bodies.

Similar services are beginning to emerge elsewhere.

A Google Ventures-backed firm is applying AI to legacy modernization end-to-end, re-engineering mainframe systems as behavioral digital twins rather than code translations.

Even the frontier model companies see the operations gap. OpenAI and Anthropic have both created deployment vehicles, with a funding of \$4B+ and \$1.5B, respectively. They primarily target middle-market and PE-backed companies – firms that lack the deep systems-integrator relationships larger enterprises rely on to put AI into production. But the signal matters: the model companies themselves are investing billions to close the gap between AI capability and operational delivery. This is not only an insurgents' game. Some incumbents are starting to move, tying their pricing to results instead of headcount, and reorganizing teams to own a process rather than just staff it. Early signs suggest traction: the firms repositioning around outcomes are winning the larger deals. The point is not that established players can't adapt; it is that adapting means changing the business model, and committing to a position in the market narrow enough to prove, not just broad enough to claim.

None of this is as settled as the momentum suggests. Outcome-based deals still hit slow buying processes, cautious legal teams, and clients who love the idea in the room but go back to the old way when it's time to sign. The shift is real but uneven, and further along in operations than elsewhere. What is not in doubt is the direction: the money now increasingly favors those who own the result over those who sell the hours, and that pull only grows as the technology gets better.

The advantage is real, but it is not yet locked. The window to move from selling effort to owning outcomes is open, for now.

4. The Advantage Hiding in Plain Sight

Here is what the most new entrants do not have: your client relationships, your knowledge of the plumbing, and your balance sheet. But here is what they may not suffer from: margin pressure on existing contracts, client demand for discounts on the old model while trying to build the new one, and board expectation about quarterly growth.

A new entrant can build the right economic model from scratch. But they start from nearly zero on the things that make an AI agent accurate in a specific client environment: the exception patterns that took years to learn, the regulatory quirks that only surface in production, the integration points where data flows break, and the institutional knowledge of how this particular client's process actually works versus how the documentation says it works.

An incumbent who has been managing a client's supply chain operations for ten years, or handling their customer operations across multiple markets for eight, or running their compliance and risk workflows for twelve, already has this knowledge.

It is embedded in the team, the runbooks, the escalation patterns, and the exception-handling logic. This is the domain intelligence that makes an AI agent reliable in production. The advantage holds only when that knowledge is treated as an institutional asset – captured in the AI layer, the process maps, and the exception libraries, where it compounds and can be redeployed across clients. Done right, the shift to an AI-led model is what can turn hard-won know-how into something the firm durably owns. That is the path to a compounding advantage, but only if the incumbent converts it before the knowledge erodes or the client decides a startup will figure it out faster.

The incumbent also has the balance sheet: the people who understand the domain, the capital to invest, and the client trust to propose a different structure. But none of this may convert automatically. This could involve convincing clients to co-invest in the new model while the old one is still running, convincing internal stakeholders that short-term margin compression is the price of long-term relevance, and doing all of this from a position of strength before the market forces it from a position of weakness. It is a narrow window in which incumbents can act from choice rather than necessity – and those who move first will not just defend their position, they will define it.

The raw material for a compounding advantage exists inside most incumbents today. And every quarter it sits unconverted, a new entrant builds its own domain depth from scratch.

The incumbents who move first will not do so by transforming everything at once. They will do it by converting their deepest advantage – one process, one client into proof.

The first move does not need to be an enterprise-wide transformation. It is one process and one client. Pick the process where you have the deepest domain expertise and the most measurable outcome. Build the operating model around owning it end-to-end – the commercial structure, the SLA, the operating team, and the AI layer. Prove it compounds. Then go to the next client and say: here is what happened. That proof is your runway.

This is simple to describe and hard to do; the difficulty is not the technology but the organizational change it forces. Doing it well means getting four things right from the start. One, structuring the commercial model around the outcome, so the client is buying a result and you are rewarded for delivering it. Two, building a dedicated team that owns the process and is measured on what it delivers, not on utilization. Three, putting the domain knowledge into the AI layer – the runbooks, the exception patterns, the process maps – so the advantage lives in the system and compounds with every cycle, rather than residing only with the people. And four, securing internal backing before you start, so the first engagement has the internal backing to trade short-term margin for long-term position.

5. The Window is Compressing

A large share of the enterprise AI contracts signed during the 2023-24 hype cycle come up for renewal by late 2026, creating the largest reallocation moment in enterprise technology in a decade. At the same time, the EU AI Act's obligations for high-risk AI systems took effect in August 2026 - meaning companies that deploy AI in areas like hiring, credit, healthcare, or critical operations will have to prove those systems are governed, transparent, and accountable, or face steep penalties. That pushes buyers even harder toward partners who can stand behind outcomes and demonstrate control, not just hand over the technology.

Enterprises that have fully integrated AI with an accountable operating partner are showing measurably stronger results than those still piloting.

For leadership teams, the diagnostic is straightforward:

- Where does AI progress make our current business model stronger?
- Where is our differentiation, and where do we have the right to claim it?
- Which parts of our portfolio are exposed to compression?
- Where do we have the domain depth, client trust, and operational knowledge to own the outcome, not just enable it?
- What is the one process where we can prove the compounding model – and what must change to make it work?

The Bottom Line

AI improvement is not a rising tide. It is a widening wedge. On one side: companies that compound with every upgrade. On the other: companies that get repriced by it. Clients are not asking for better tools; they are asking for owned results. Most providers have the raw material to deliver them: the domain depth, the client intimacy, and the operational knowledge no startup can replicate from scratch. But converting the raw material into a compounding advantage requires runway, conviction, and a willingness to build the new model before the old one stops working. The window to do this from a position of strength is narrow. The companies that use it will set the terms for the next decade. The ones that wait will negotiate from weakness – or not at all.

About the analysis

This perspective draws on BCG's analysis of 40+ players and 100+ positioning claims; 70+ CXO interviews across banking, insurance, healthcare, retail, and manufacturing; and market data from Gartner (IT Spending Forecast, April 2026), IDC, a16z, Stanford HAI AI Index Report 2025, and Sequoia Capital. Enterprise survey data on AI failure rates and CEO risk perception from HCLTech/Omdia AI Impact Imperatives 2026 (n=467, \$1B+ enterprises). Company data on Palantir, OpenAI, Anthropic, and select AI-native services players from public disclosures, earnings calls, funding announcements, and product launches. The supporting Market Signal Report, competitive analysis, and capability taxonomy are available as a separate research companion.

– BCG Tech Capital, Growth Tech, and IT Services Practices

About the Authors



Rajiv Gupta

Managing Director and Senior Partner
New Delhi
Gupta.Rajiv@bcg.com



Akash Bhatia

Managing Director and Senior Partner
San Francisco
Bhatia.Akash@bcg.com



Mandeep Kohli

Managing Director and Partner
New Delhi
Kohli.Mandeep@bcg.com



Karan Chadha

Managing Director and Partner
New Delhi
Chadha.Karan@bcg.com



Jade Leet

Brighthouse Managing Director
Atlanta
Leet.Jade@bcgfed.com



Cara Castellana Kreisman

Brighthouse Senior Strategy Director
Atlanta
CastellanaKreisman.Cara@bcgbrighthouse.com

Acknowledgments

We extend our gratitude to Augustya Mudgal, Ananya Kalra from BCG India, Nopur Nopur from BCG India Marketing Team. We would also like to thank Saroj Singh, Ratna Soni, Alivya Saha, and Soumya Garg for their contributions to the editing, design, and production of this report.



About Boston Consulting Group

Boston Consulting Group bridges the gap between ambition and outcomes for the world's leading companies and organizations. We are built for this era of unprecedented change — bringing strategic clarity rooted in over 60 years of deep domain knowledge, combined with applied AI shaped by our practitioners. BCG works shoulder-to-shoulder with CEOs across industries and geographies to deliver transformative impact at scale: stronger returns, transferred capabilities, and change that sticks. For more information, visit bcg.com.

For information or permission to reprint, please contact BCG at permissions@bcg.com. To find the latest BCG content and register to receive e-alerts on this topic or others, please visit bcg.com. Follow Boston Consulting Group on [LinkedIn](#), [Instagram](#), [Facebook](#), and [X](#).

