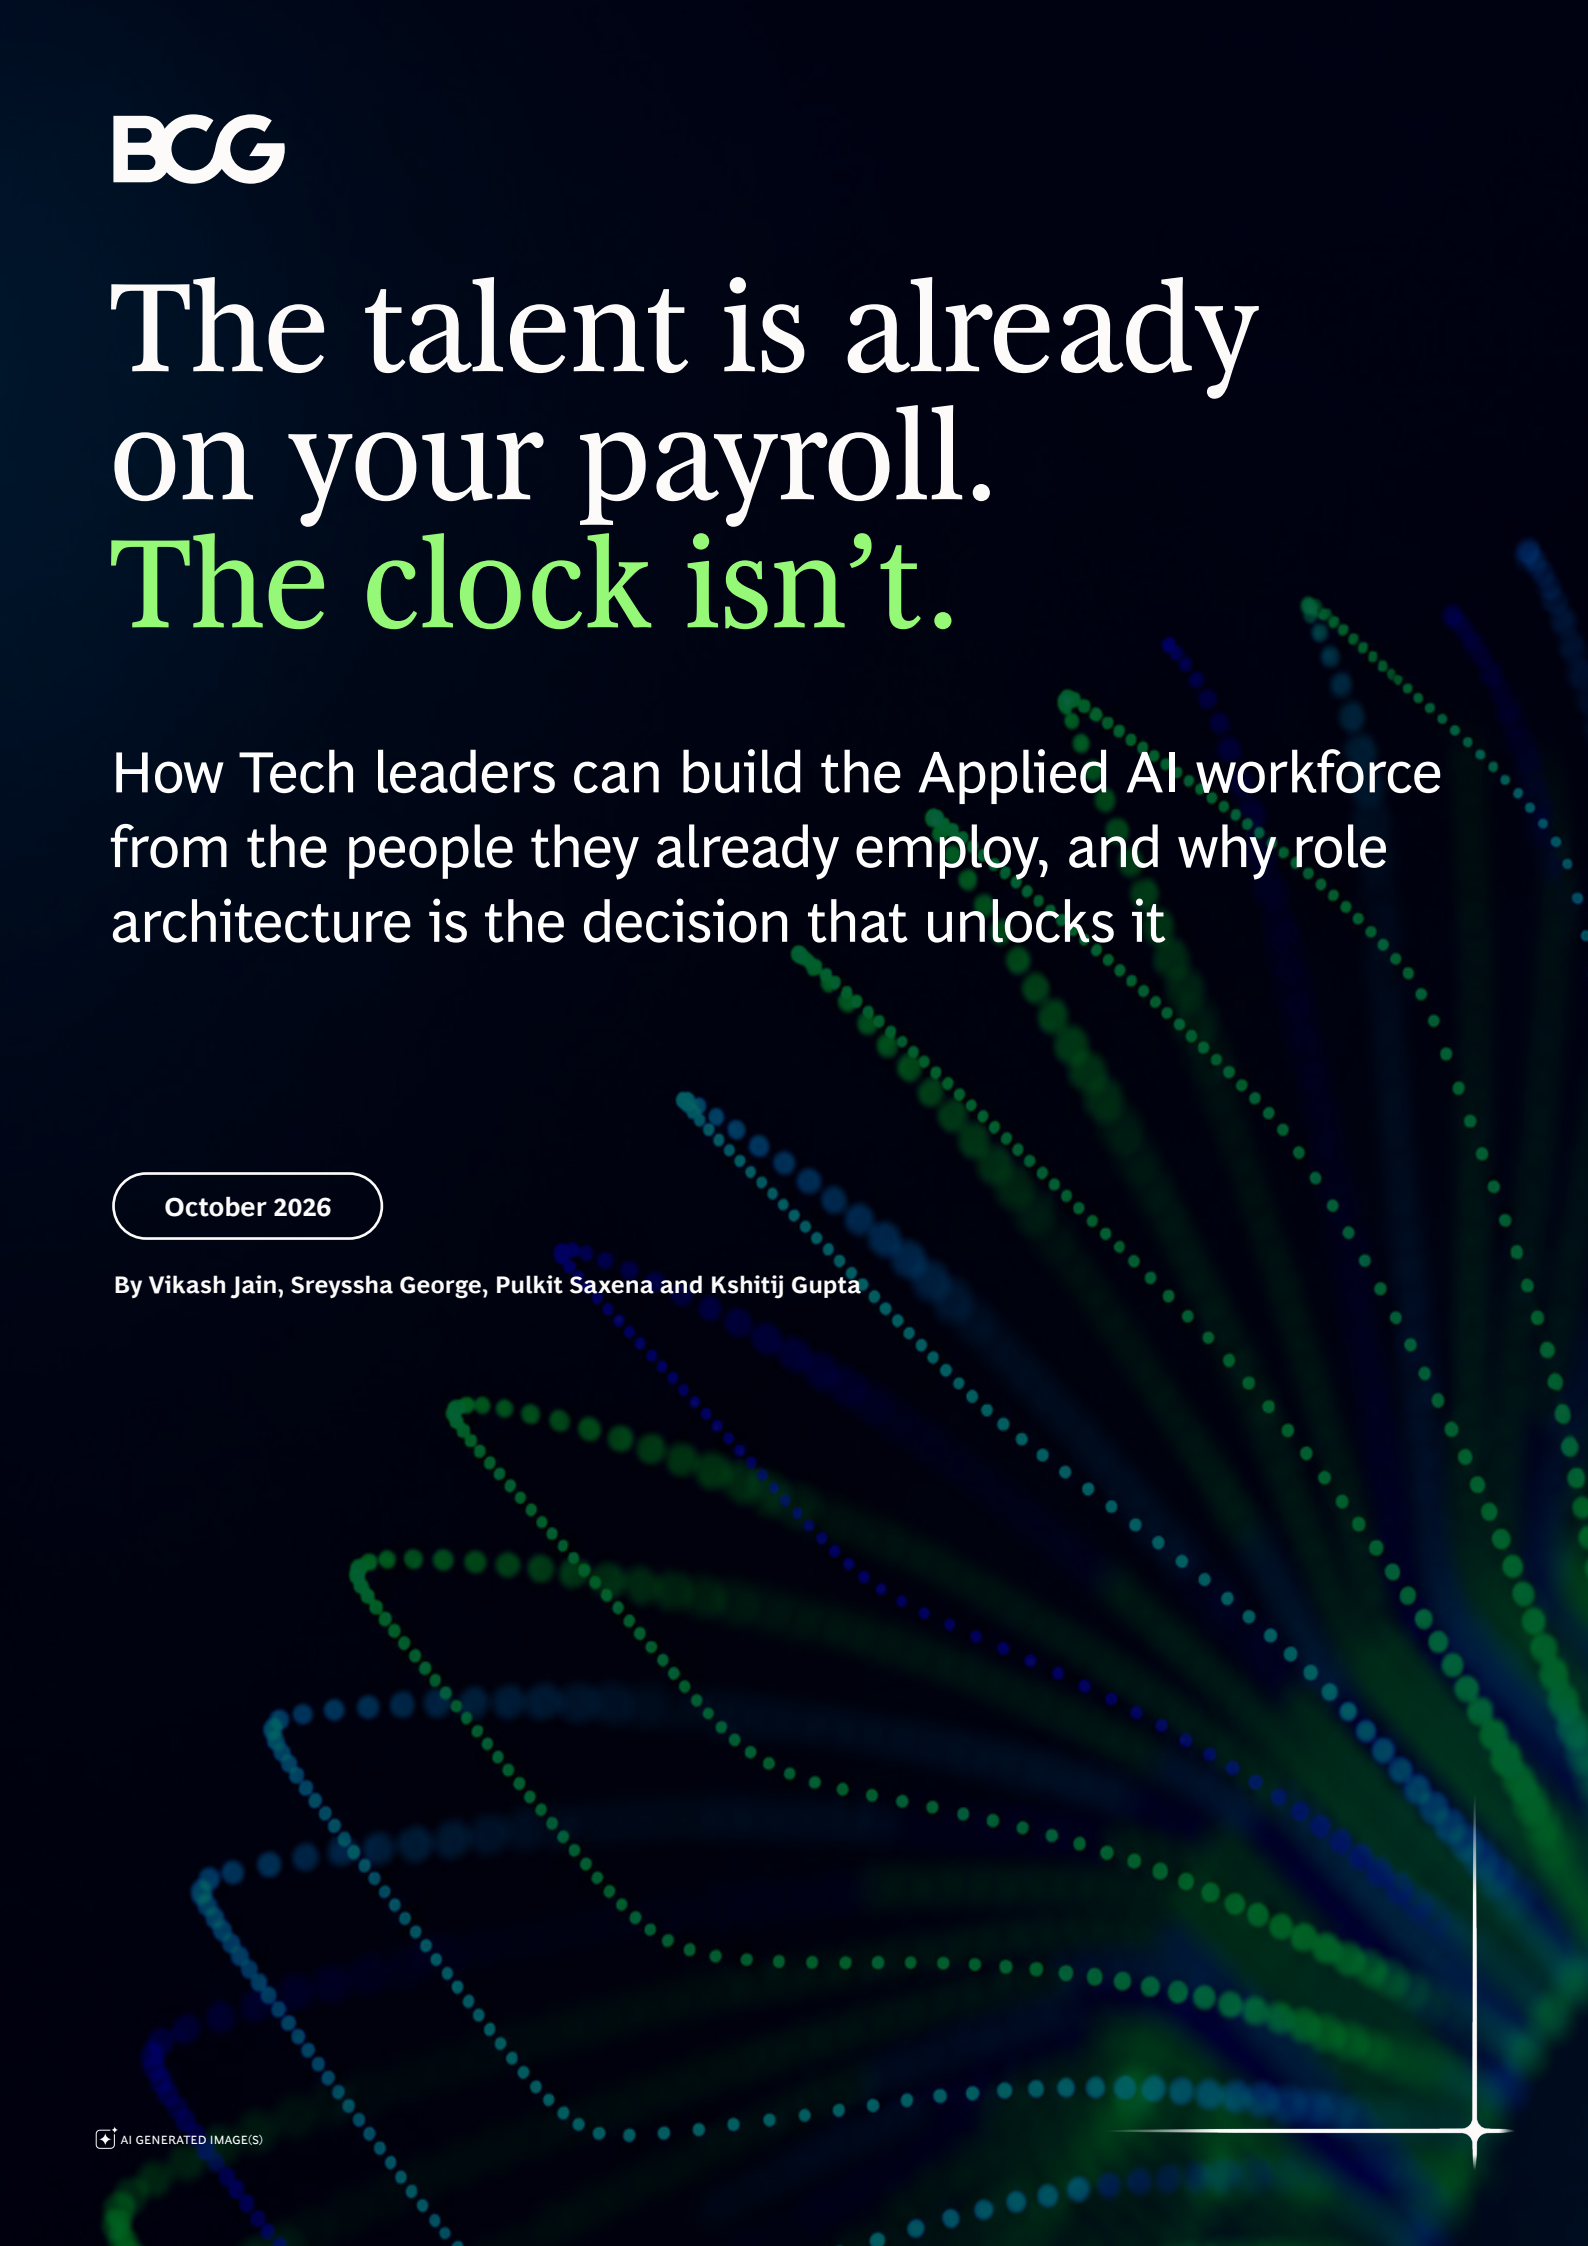
The talent is already on your payroll. The clock isn't.

How Tech leaders can build the Applied AI workforce from the people they already employ, and why role architecture is the decision that unlocks it

October 2026

By Vikash Jain, Sreyssha George, Pulkit Saxena and Kshitij Gupta



AI GENERATED IMAGE(S)



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.

Table of Contents

1 Applied AI presents a $100-200B opportunity; but tech SIs have only a fraction of the talent needed to capture it

2 Over 2,500 AI-tagged postings highlight 40 Applied AI roles – each calling for distinct talent decisions

3 AI natives place ~45% of their Applied AI talent in forward-deployed roles; Tech SIs place under 5%

4 Every one of the 40 roles has an internal lineage, and the gaps are specific enough to close

5 Role architecture requires detailing across four key dimensions

6 One Applied AI role could open into multiple career tracks, and that structure outlasts any pay band

7 Five questions to answer this quarter

Today, Applied AI is sold on a delivery model most Tech SIs have not yet built

Applied AI work is being sold and delivered on different economics from traditional IT services. Applied AI here means building, deploying and running AI-native and agentic solutions in production, for clients. Infused AI covers the use of AI tools to speed up work that already existed. Across AI-native firms, close to half of Applied AI talent sits in forward-deployed roles: senior, client-embedded, accountable end to end from discovery through delivery. At Tech SIs, forward-deployed roles account for under 5% of the Applied AI workforce, and Applied AI itself is only 10-15% of their AI-tagged talent. Two delivery models now compete for the same client budget on different unit economics.

The capability behind the embedded model can be built largely from the workforce already in place, with external hiring kept selective: the scarce anchor roles, forward-deployed first, and top-up hiring for scale-driver roles while internal cohorts build. Applied AI hiring across the market resolves into 40 distinct roles, and each one has a direct lineage to a role a Tech SI already employs: a forward-deployed engineer evolves from a data scientist or a full-stack developer, an AI Ops engineer from a DevOps or Site Reliability Engineer (SRE) profile, a semantic engineer from a data engineer.

The decision that turns those lineages into supply is role architecture: skill definitions, proficiency gates, bands and pay logic, so that a data engineer can become an AI data engineer with a title, a level and a band attached. Firms that build it create their own supply. Firms that hire against the new titles pay up to roughly 35% over the comparable traditional role for a capability much of their workforce is already close to.

This report sets out where the opportunity sits, which roles carry it, how far the firm's existing people are from them, and the architecture and career decisions that separate the two outcomes.

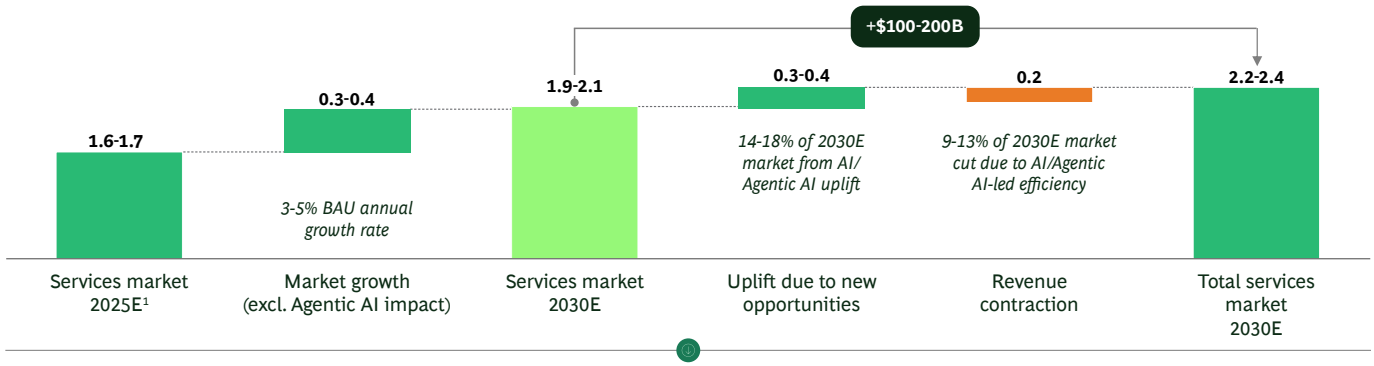


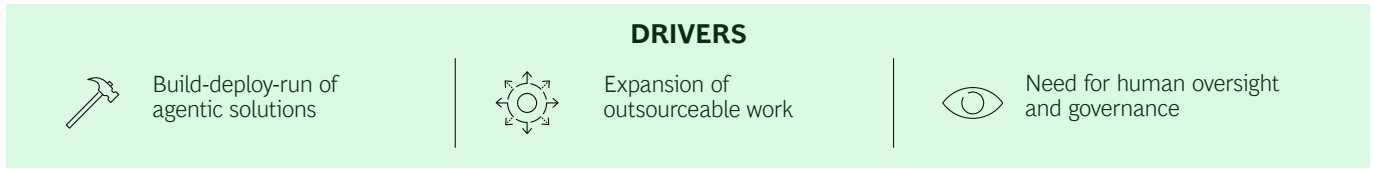
1. Applied AI presents a $100-200B opportunity; but tech SIs have only a fraction of the talent needed to capture it

The services market is projected to grow from roughly $1.6-1.7T in 2025 to $2.2-2.4T by 2030 (Exhibit 1). AI-led services account for $300-400B of that growth, driven by three forces: (1) building, deploying and running agentic solutions rather than staffing them; (2) absorbing a wider slice of work that becomes newly outsourceable once agentic tooling matures; and (3) supplying the human oversight and governance layer that clients may not delegate to a model alone. The net is a $100-200B Applied AI opportunity, available to the firms that can deploy the new capabilities at scale.

EXHIBIT 1

AI-led services expand the market while reshaping where value is created





Source: BCG report - Agentic AI Opportunity for Tech Service Providers, February 2026
[1]Services market comprises five segments: Application Implementation and Managed Services, Consulting, Infrastructure Implementation and Managed Services, Business Process Services, and Infrastructure as a Service (IaaS) .

Few firms are set up to capture that today. The constraint is the composition of the AI workforce. Across 50,000+ AI-tagged professional profiles, Applied AI work – building and running agentic systems rather than using AI to accelerate existing tasks – accounts for 90-95% of the AI workforce at AI-native firms and 35-45% at digital engineering specialists. At Tech SIs the figure is 10-15%. The remaining 85-90% at Tech SIs is Infused AI: data analysts, functional consultants and engineers applying AI to work that already existed.

10-15% of the AI-tagged workforce at Tech SIs does Applied AI work today, against 90-95% at AI-native firms

Tech SIs appear to have hired their way to much of today's 10-15%, paying a premium for each specialist. That works for pilots. Production delivery needs full teams at every level. The durable answer is to build the engine internally: define Applied AI roles and a career architecture that draw on the existing 85-90%, so the firm creates its own supply.

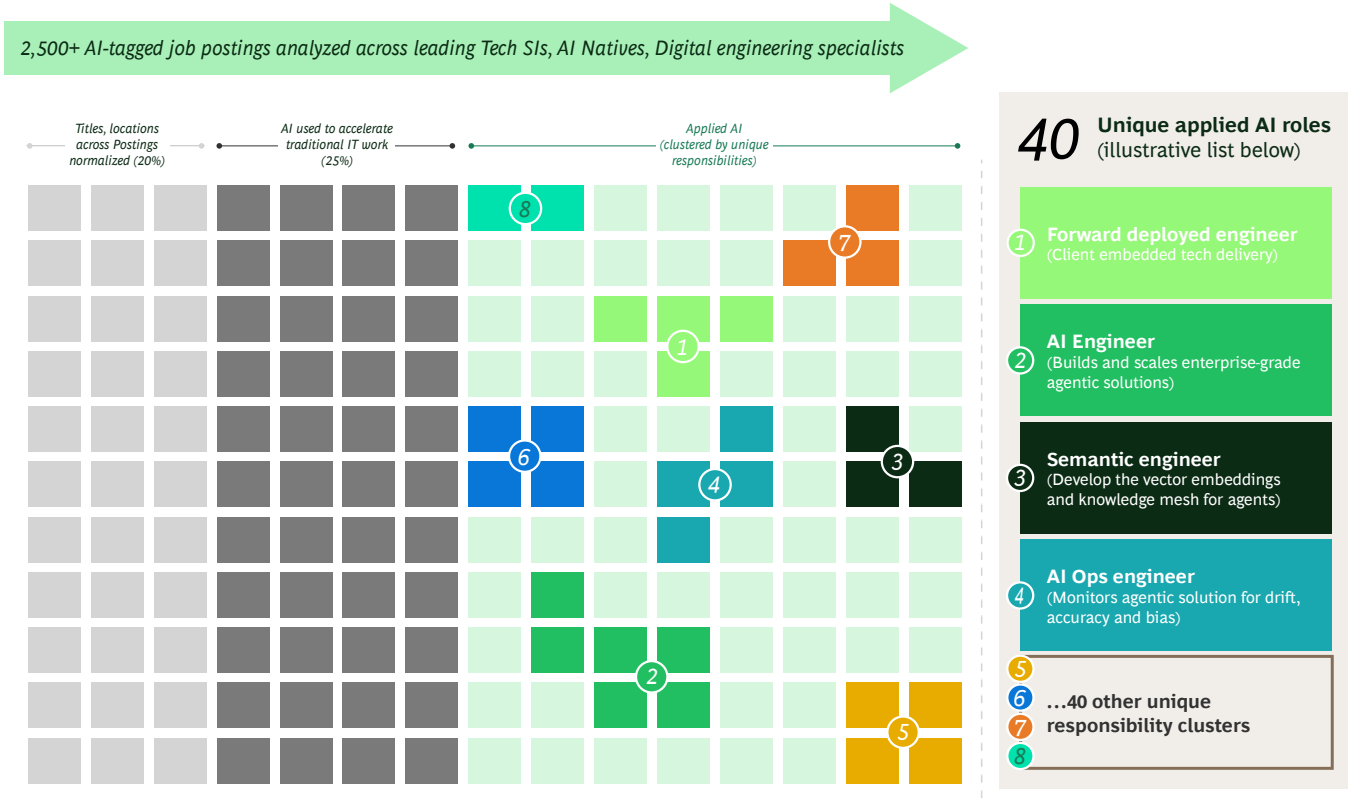


2. Over 2,500 AI-tagged postings highlight 40 Applied AI roles – each calling for distinct talent decisions

More than 2,500 AI-tagged open positions were screened over six months across AI-native firms, digital engineering specialists and Tech SIs. On deeper analysis and excluding postings that simply use AI to accelerate existing work, the remainder cluster into 40 distinct Applied AI roles (Exhibit 2).

EXHIBIT 2

Market evidence distilled from 2,500+ AI-tagged postings into 40 distinct Applied AI roles.



Source: 2500+ Job postings and 50K+ profiles of people with AI-keyword extracted from company job boards and third-party job portals across 15+ firms categorized into Tech SIs, AI natives and Digital engineering and analytics firms; Detailed postings assessed for relevance to applied AI work and the responsibilities clustered uniquely to synthesize into 40 unique applied AI roles

Forty is a workable number. It is small enough to take a deliberate decision on each role, and the roles do not all warrant the same decision. Plotted against market signal and the number of such profiles currently sitting inside Tech SIs, they sort into four groups, each with its own action.

- **Scarce anchors – hire selectively, pay the premium.** Strong market signal, limited internal population: forward-deployed strategist, forward-deployed engineer, digital workforce manager. A small number of external hires here gives the capability a visible role model.
- **Scale drivers – hire actively and build cohorts.** Strong signal and a growing internal base: AI engineer, AI architect. These carry the delivery model and warrant both active hiring and the largest upskilling cohorts.
- **Delivery core – staff to retain and grow.** Lower market noise, significant internal population, heavy delivery load: AI data engineer, AI delivery owner, full-stack developer. Retention risk here is routinely underrated, because these are the roles an upskilling programme touches first.
- **Early signals – assign an owner and watch.** Emerging roles with thin signal and thin populations: AI experience designer, AI FinOps. These warrant a named owner and a quarterly review.

Four groups give four decisions – buy, build at scale, retain, or watch – across a 40-row list that would otherwise carry a single undifferentiated plan.





3. AI natives place ~45% of their Applied AI talent in forward-deployed roles; Tech SIs place under 5%

Applied AI delivery models are still taking shape, and they differ by the kind of firm doing the work. AI natives are largely platform businesses whose work is deploying their own platform inside a client. They therefore staff heavily with forward-deployed engineers and strategists, together around 45% of their Applied AI workforce. This is the ***Fully Embedded*** model, in which small, senior, client-embedded teams own discovery, production-ready build and the outcome.

Tech SIs are still working out who runs the pilot and who runs the scale-up. Most operate a ***Hybrid Spearhead***, in which a lean client-facing team owns scope and the relationship while scaled engineering hubs execute delivery. Their Applied AI hiring spreads across full-stack, AI/LLM, architecture and data engineering roles, and forward-deployed roles account for under 5%. Tech SIs that own a platform tend to sit closer to the AI-native mix (Exhibit 3).

EXHIBIT 3

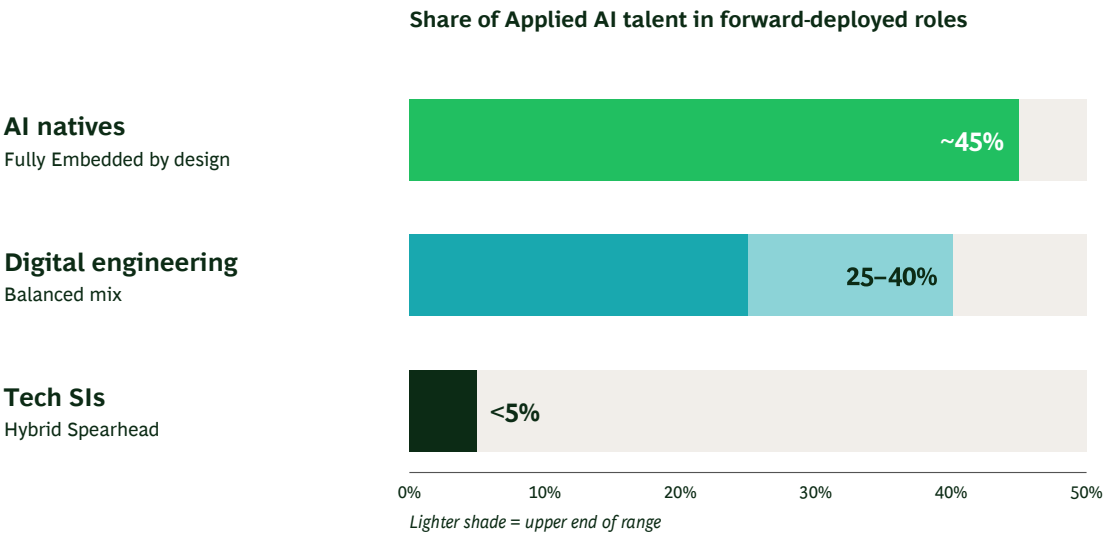
Delivery model choice drives the Applied AI talent

Fully Embedded

Small, senior teams sit inside the client and own discovery, production-ready build and the outcome. Staffed with forward-deployed engineers and strategists.

Hybrid Spearhead

A lean client-facing team owns scope and the relationship; scaled engineering hubs run delivery. Forward-deployed talent is used selectively.



More than 75% of firms in each group are hiring for their lead model: embedded for AI natives, hybrid for digital engineering and Tech SIs. The pattern holds across the group and does not rest on a few outliers.

Source: BCG analysis of 2,500+ AI-tagged job postings and 50K+ AI-tagged profiles across 15+ firms (AI natives, digital engineering specialists, Tech SIs).

Both models win work, but they price it differently. A five-person embedded team with agentic platforms / accelerators that owns an outcome and a large time-and-materials team working a statement of work are not interchangeable in a client's evaluation.

Two decisions follow, and both should be made deliberately:

- **Pricing must reflect outcome ownership & delivery scale.** Once a client has seen an embedded team own an outcome, that becomes the reference point. To meet these expectations, hybrid spearheads must shift toward value- and outcome-based pricing, backed by delivery at scale. For instance, a finance-operations automation programme priced on time and materials could instead be priced per transaction processed. Without outcome-based pricing, a forward-deployed engineer is a high-cost resource sold on a rate card built for hours, so role design and the commercial model have to be settled together.
- **Applied AI talent needs clear roles, ownership and progression.** AI-native firms offer the senior, ownership-heavy roles that AI professionals seek. Hybrid spearheads can offer similar ownership alongside the breadth and scale of enterprise work. Firms that want to win must give in-house talent clearly defined roles, explicit responsibilities and visible career progression.



4. Every one of the 40 roles has an internal lineage, and the gaps are specific enough to close

The adjacencies hold up when roles are mapped against a real Tech services repository of 200 roles across 2,000 skill clusters and 3,000 skills (before deduping closely related skills). Each role builds on capabilities already in your workforce, with targeted upskilling needed to bridge the gap: a data scientist or full-stack developer can develop into a forward-deployed engineer; a DevOps or Site Reliability Engineer (SRE) into an AI Ops engineer; and a data engineer into a semantic engineer.

A single lens can mislead here. Job titles, skills and prior career paths each reveal only part of the picture. Firms need a holistic view of all three to assess readiness for Applied AI roles and identify the upskilling needed, without overestimating or underestimating their people's capabilities.

New skill requirements are emerging for Applied AI roles, though the mix and depth vary by role. Key skills to assess against each role's requirements include the following (non-exhaustive):

- **The retrieval stack.** RAG engineering and evaluation, chunking and reranking, vector databases and vector search – the layer that grounds a model in enterprise data, and the most consistent gap in the set.
- **Evaluation and safety.** Evaluation frameworks and harnesses, bias detection, explainability, guardrails.
- **Agentic orchestration.** Multi-agent systems, routing and model selection, tool, memory and planning.
- **Integration into enterprise platforms.** Wiring models into the systems a client already runs, now treated as table stakes.

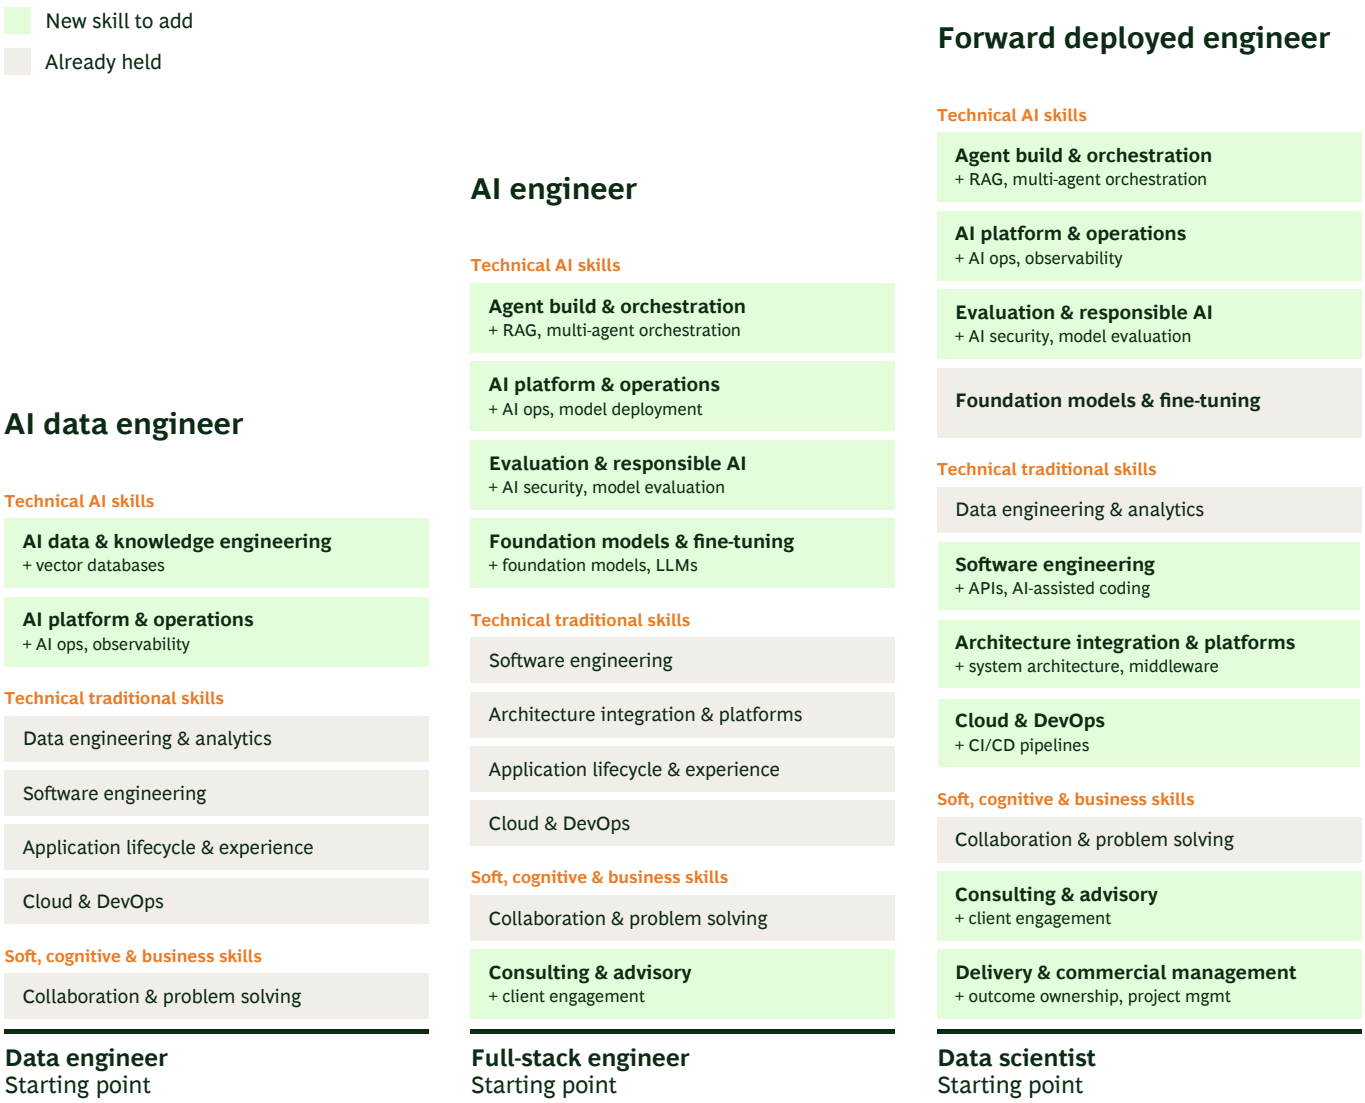
The transitions vary sharply in difficulty (Exhibit 4A). A data engineer moving into an AI data engineer role is a contained build-out, concentrated in vector databases and AI platform operations. A full-stack engineer becoming an AI engineer needs a broader set of AI-specific capability. A data scientist moving to forward-deployed engineer is the most demanding, adding client advisory, commercial ownership and end-to-end delivery.

3x more upskilling effort to move a data scientist to a forward-deployed engineer than to move a data engineer to an AI data engineer

EXHIBIT 4A

Building Applied AI capabilities from the closest IT role requires varying degrees of skill change

AI roles stack new skills on a base of existing ones: the data scientist path needs the tallest add-on



Source: BCG skills and proficiency analysis. Skill requirements mapped per role across a three-family taxonomy and classified as must-have or good-to-have from job description analysis across peer firms and expert interviews.

Multiple roles feed the same destination from different capability bases, so the journey is mapped per person (Exhibit 4B). Take the forward-deployed engineer:

- A data scientist typically arrives strong in foundation models and data and analytics. They need to build engineering and platform depth (agent orchestration, AI platform operations, cloud and DevOps) as well as client advisory and commercial ownership.
- A full-stack developer brings software engineering, application lifecycle and cloud skills. They need most of the AI-specific layer: agent build, foundation models, evaluation and AI platform operations.
- A data engineer brings data, integration and platform skills. They need AI-specific build capability and the full client-facing layer.

Conversion has a delivery cost that the plan has to carry. Sequence cohorts against utilisation, and attach a cost and a time to productivity to each route in Exhibit 4B.

Some adjacent profiles will not convert. The plan needs a clear position on the Infused AI workforce: who moves into Applied AI roles, who is redeployed into adjacent work, and where attrition is acceptable.

EXHIBIT 4B

Three routes into the forward-deployed engineer role, each with a different skill gap to close

Traditional IT roles offer viable transition paths to the FDE profile; mapping the skill gaps define individual upskilling journeys[1]



Source: BCG skills and proficiency analysis. Skill requirements mapped per role across a three-family, 15 -group taxonomy and classified as must-have or good-to-have; job description analysis across peer firms and expert interviews. Ordering within each block reflects relative scarcity and is directional
1. Role profiles derived from a JD-to-JD comparison of published postings; actual capability varies by individual, organisation and engagement context



5. Role architecture requires detailing across four key dimensions

When digital transformation took hold, the firms that scaled fastest built supply by defining new roles, such as UX designers, cloud architects, and DevOps and full-stack engineers. They built job families around those roles and moved people across from existing delivery roles. Applied AI calls for the same move, with internal supply as the goal. The building blocks are familiar: a skill taxonomy, proficiency levels, bands and a job family. What matters is where Applied AI forces each one to work differently.

1. **A skill taxonomy granular enough to assess against.** Our analysis of 2,500+ AI-tagged job openings, combined with BCG's proprietary skill repository, identified roughly 2,000 unique skills relevant to Applied AI work. These are organised into 30 skill groups across three families: Technical AI; Technical Foundational; and Soft, Cognitive and Business. This granularity makes "qualified for this role" a measurable statement.
2. **Proficiency levels per skill, with evidence attached.** Each sub-group carries four proficiency levels, from beginner to expert, with clear evidence required at each level. For prompting, for instance, a beginner can write a clear prompt for a routine task; an expert can design and refine prompting strategies for complex tasks, systematically test outputs and resolve failures.
3. **Bands that hold against the existing structure.** Applied AI roles map onto existing bands with differentiated, market-facing titles, while the band economics underneath stay intact. The premium could be paid as a hot-skill incentive or a benchmarked base on transfer or selective re-banding, which avoids a pay-parity conflict with traditional delivery roles.
4. **A dedicated job family.** Firms that have moved fastest created new job families to hold the distinct skill clusters. An Applied AI role placed inside an existing delivery family inherits that family's progression logic and pay band along with it.

Together these produce what a training programme on its own cannot: a named destination role, a measurable gate, a band and a title.

6. One Applied AI role could open into multiple career tracks, and that structure outlasts any pay band

Applied AI roles command a premium today, falling into three bands against the closest comparable traditional IT role: roughly 5-15%, 15-25% and 25-35%, with the top band holding the forward-deployed and research-oriented roles. A firm hiring externally into those roles pays that premium on every hire.

Premiums are a moving target. As Applied AI skills diffuse through the market, the scarcity behind today's top band compresses and roles migrate down the tiers. Re-banding salaries across a large delivery organisation every time a skill commoditises is slow, expensive, and sets a precedent that is difficult to unwind. Career structure is the durable lever: it holds its value as the skill premium moves.

25-35%: Premium forward-deployed roles command over the closest traditional IT role

Three things make the lattice work:

- **Five destinations out of one role.** A single ladder gives an AI engineer one destination and a queue for it. The role opens into five tracks: individual contributor and architecture; forward deployment; managerial; research; and product and platform. An engineer with no interest in managing people advances into architecture or frontier research; one who wants client ownership moves into forward deployment. The track becomes the person's choice, and the firm can offer it immediately.
- **Movement sideways as well as up.** Vertical progression inside a track, horizontal movement between tracks at the same band. A forward-deployed engineer who wants to go deeper technically moves into architecture without stepping down a level, so a plateau in one track stops being a reason to leave.
- **Levels that reward demonstrated capability**. Progression is tied to the proficiency gates defined in the role architecture. Consistent performers can advance roughly a year ahead of standard band tenure at each level, and can reach leadership several years faster than they would on a time-served ladder. Band tenure becomes a default that the gates allow a firm to override defensibly, so a strong engineer is not waiting on a calendar.

The lattice compounds: each lateral move creates an internal precedent and someone who can guide the next cohort, and firms that publish their tracks use them as a recruiting argument. Pay and career structure are both tested the day a competitor calls. A number can be matched once; five visible tracks and published progression gates give a person a reason to stay that a competing offer has to beat on more than salary.



7. Five questions to answer this quarter

Role architecture, the 40 roles and the career lattice all come down to decisions a firm can take now. Five questions show where yours stands. Each question has a clock attached.

An honest “no” is a useful starting point, and the cost of leaving it unresolved rises with every quarter of delay.

Have you **prioritised the agentic use cases** you intend to build, or is the talent plan still upstream of a business decision nobody has taken yet?

01

Is there **one role architecture** for Applied AI across the firm – defined roles, a shared skill taxonomy, proficiency gates and bands – or is each business unit inventing its own titles and pay bands and quietly bidding against the others for the same people?

02

For the 40 roles, have you taken a **deliberate build, buy, retain or watch decision,** or is the default still to post the job and see what the market provides?

03

Have you **assessed your workforce against the role profiles individually, using more than one adjacency view,** or are you assuming your people are further along, or further behind, than they actually are?

04

Would your **compensation bands and career paths withstand a competing offer from an AI-native firm** today? If you are uncertain, that uncertainty is the answer.

05

Capability takes years to build; the role architecture that everything else builds on is worth settling first.

The clock is already ticking. The $100-200B Applied AI opportunity builds through to 2030, and the firms best placed to capture it are likely to be those whose people have had time to work on real engagements by then. Roles need to be defined, people need experience on real work, and the organisation needs room to adapt around them. Experience is the slowest of these to build, so firms that settle their role architecture early give their people the longest runway.

Start with the roles. Choose the ones that fit your purpose, whether that is an AI team to automate your own workflows or Applied AI delivery for clients, and give each a band, a title and a proficiency gate before the next appraisal cycle. Hiring, upskilling and pay follow from there. Every quarter without defined roles is a quarter of experience your people do not get.



About the Authors



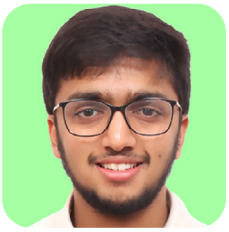
Vikash Jain
Managing Director and Senior Partner
Jain.Vikash@bcg.com



Sreyssha George
Managing Director and Partner
George.Sreyssha@bcg.com



Pulkit Saxena
Partner
Saxena.Pulkit@bcg.com



Kshitij Gupta
Project Leader
Gupta2.Kshitij@bcg.com

Acknowledgments

We also appreciate the efforts of all stakeholders and team members who provided resources, expertise and support at various stages of the process. We would like to thank Rahul Bhadauria and Jyotika Malhotra for their contributions to the research, analysis and writing of this report.

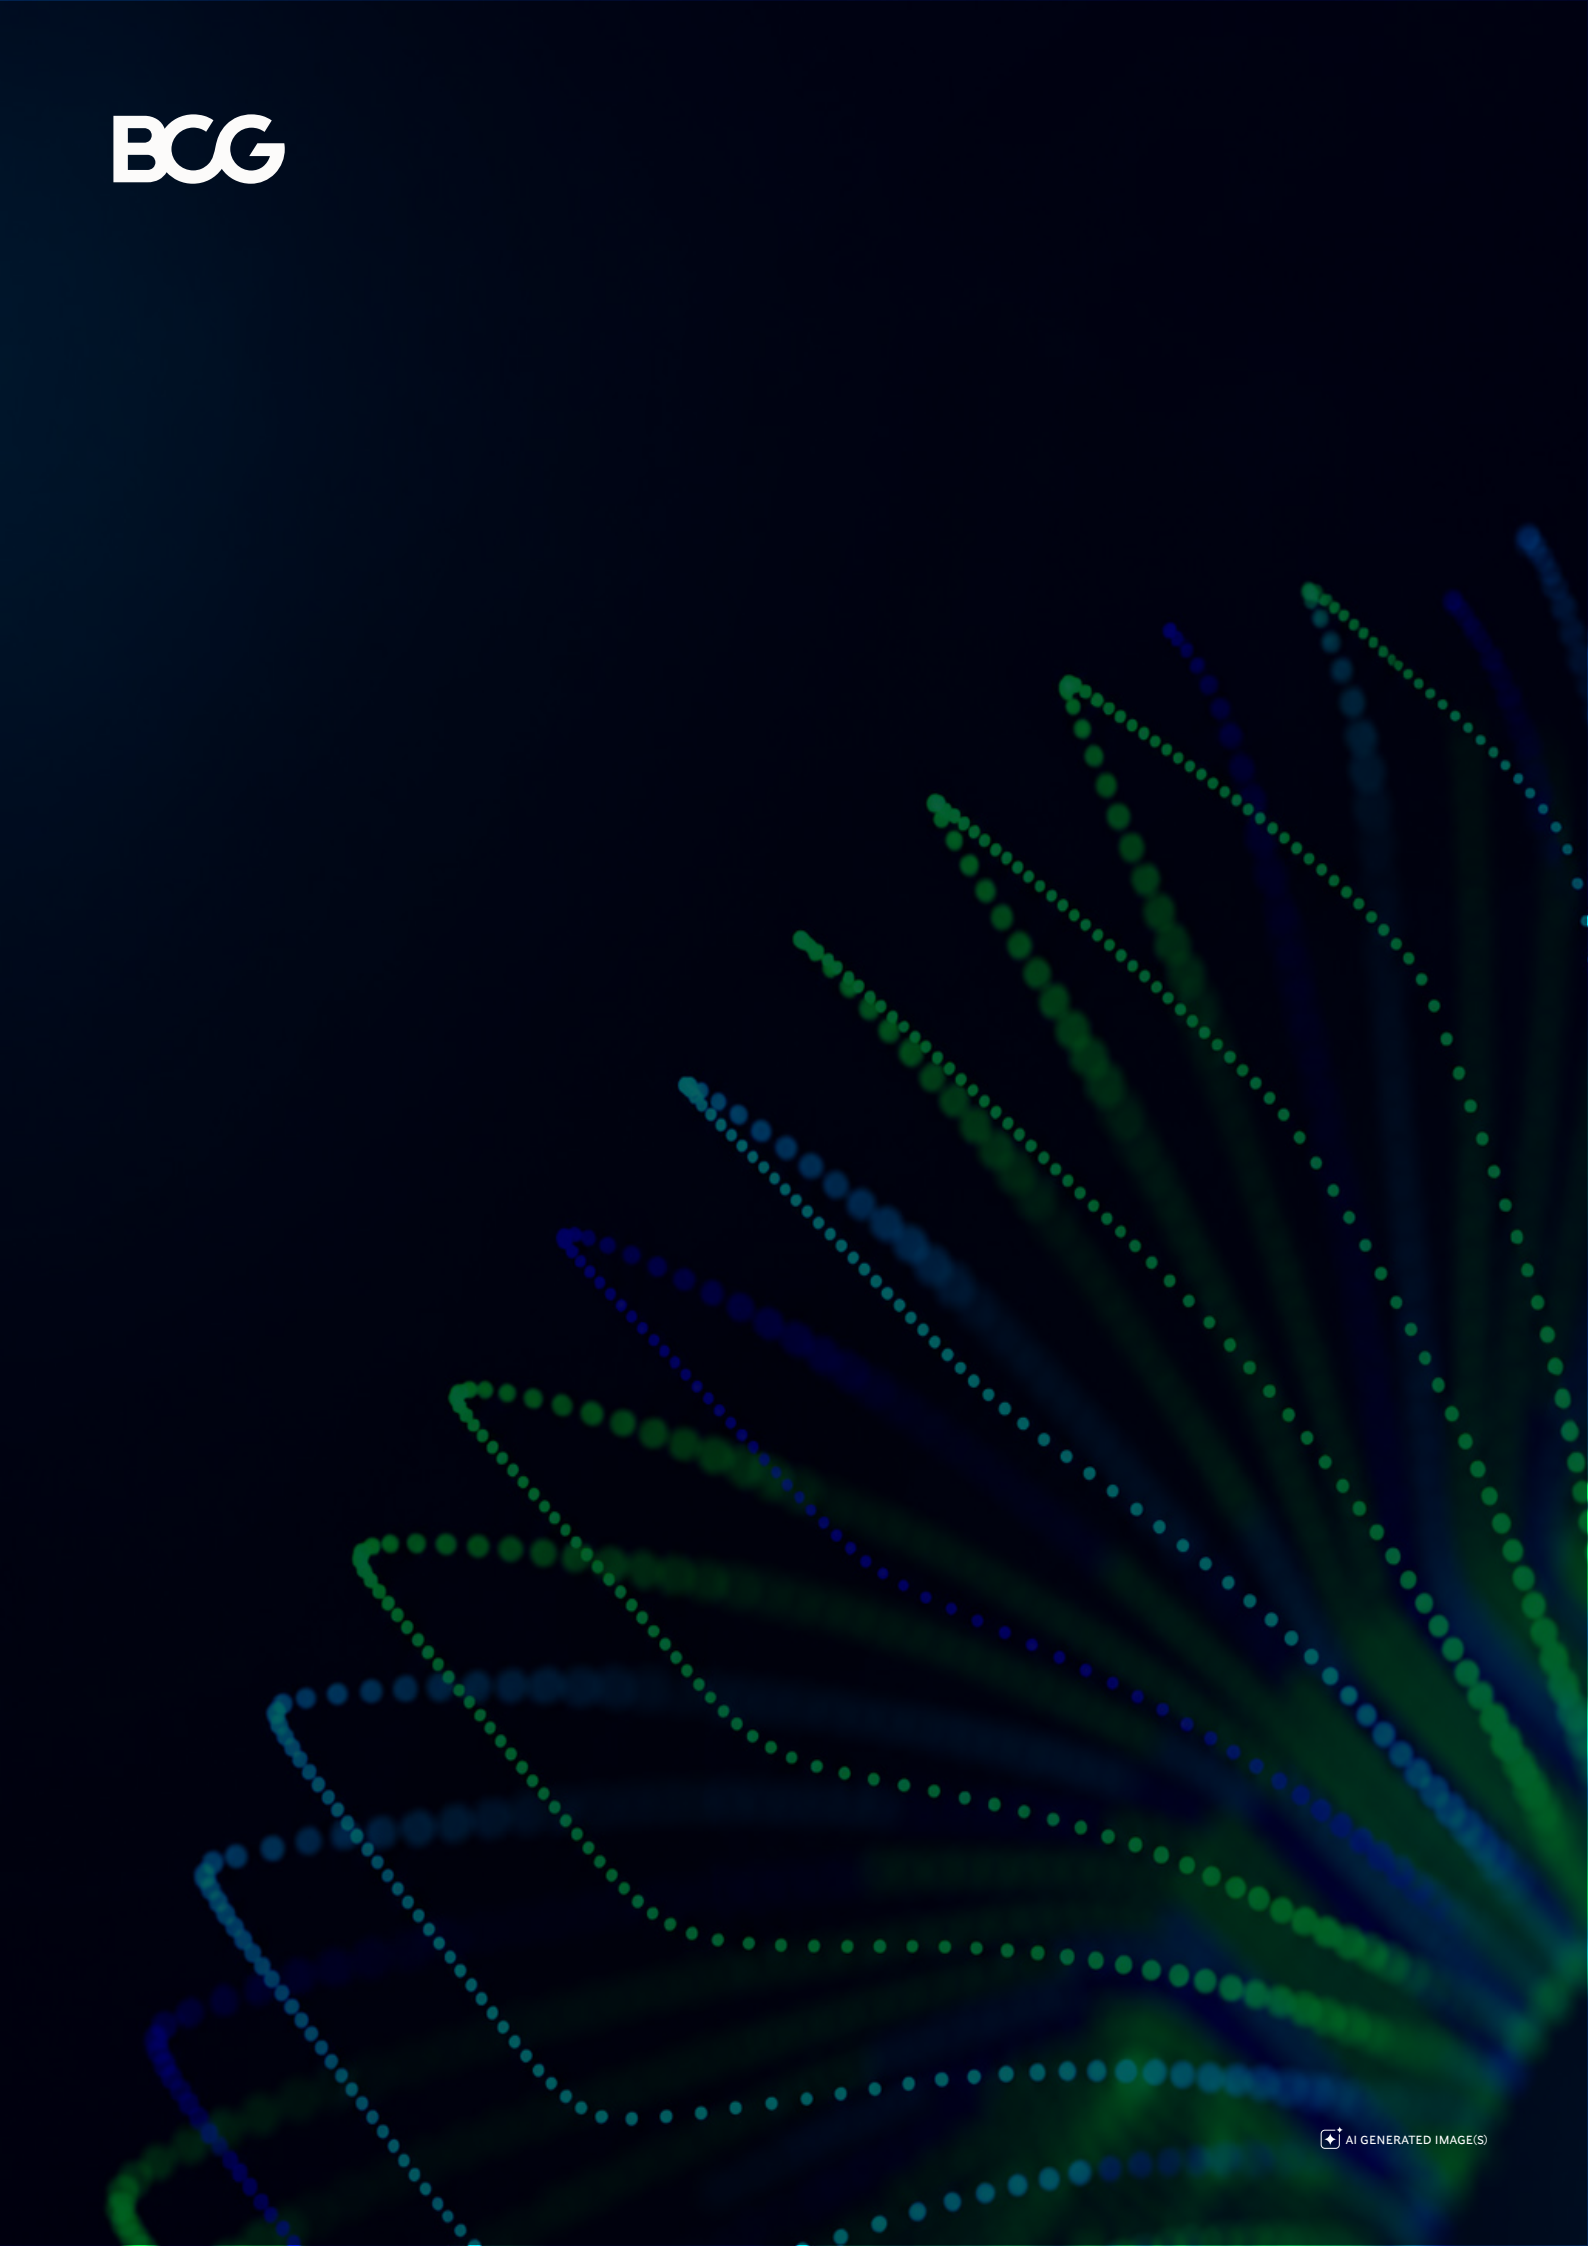
We are thankful to Nopur for her support in managing the marketing and communications process.

We would also like to thank Venkatesan Natarajan, Noyonika Mukherjee, Honesh Rai Pareek and Saurabh Tikekar for their contributions to the editing, design and production of this report.

This publication is made available solely for your personal, informational, and non-commercial use. You may download and retain a copy for such purposes. Except with the prior written permission of Boston Consulting Group (BCG), you may not use, reproduce, republish, distribute, sell, license, modify, create derivative works from, or otherwise exploit this publication or any portion of its content for any commercial purpose. Without limiting the foregoing, the publication and its content may not be used for the commercial development, training, testing, evaluation, or improvement of artificial intelligence, machine learning, or other automated systems, or incorporated into any commercial database, dataset, product, or service, without BCG's prior written permission. All rights not expressly granted are reserved by BCG.

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.

© Boston Consulting Group 2026. All rights reserved. X/26


AI GENERATED IMAGE(S)