Beyond AI Islands

Why Organizations Must Unify Disconnected Solutions to Decode The Art of The Possible

April 2025

Prasanna Arikala, Vikash Jain, Cathal McCarthy, Shaleen Sinha, Amit Kumar, Nicholas Clark, Benjamin Quirt, and Cobus Greyling

BCG + kore.ai

Boston Consulting Group partners with leaders in business and society to tackle their most important challenges and capture their greatest opportunities. BCG was the pioneer in business strategy when it was founded in 1963. Today, we work closely with clients to embrace a transformational approach aimed at benefiting all stakeholders—empowering organizations to grow, build sustainable competitive advantage, and drive positive societal impact.

Our diverse, global teams bring deep industry and functional expertise and a range of perspectives that question the status quo and spark change. BCG delivers solutions through leading-edge management consulting, technology and design, and corporate and digital ventures. We work in a uniquely collaborative model across the firm and throughout all levels of the client organization, fueled by the goal of helping our clients thrive and enabling them to make the world a better place.

Kore.ai is a leading provider of Agentic AI with over a decade of experience in helping enterprises realize business value. The Company provides strong business solutions leveraging AI for workplace, process automation and customer service use cases. These are built on a comprehensive agent platform that brings together autonomous agents, sophisticated enterprise knowledge retrieval, intelligent agent orchestration and no-code/pro-code tools. Kore.ai takes an agnostic approach to model, data, cloud and applications, giving customers freedom of choice. Trusted by over 450 Global 2000 companies, Kore.ai is helping the navigation of AI. Visit Kore.ai to learn more.

Contents

Executive Summary

Enterprise AI is at an inflection point: **Three in four technology leaders fear 'silent failure' in AI**—where millions of dollars are invested in solutions that appear functional but quietly fail to deliver bona fide business value or trustworthy outcomes.

Q: What's your biggest fear about deploying AI agents across your organizations?

Most of these failures stem from enterprises optimizing small-scale wins without revisiting the operating model. Such risks will only grow with the integration of more AI solutions into enterprise workflows[1].

Disconnected AI deployments are increasingly exposing enterprises to unintended consequences such as duplicative efforts, security vulnerabilities, compliance risks, maintenance and scalability challenges. Islands of AI are popping up all over the enterprise—but without coordination between technology and business functions, significant value is lost.

Addressing these challenges demands a balanced approach, integrating the strengths of specialized AI point solutions (precision and agility) with the comprehensive benefits of centralized AI platforms (governance, reusable AI components, and enhanced scalability) to create a unified, hybrid AI ecosystem across an enterprise. Platforms enable a shared system to manage, govern, deploy, and scale AI across enterprises as opposed to point solutions which are designed for a particular use case, business unit or a function.

BCG conducted a study on the key evaluation criteria for AI solutions, surveying over 75 CIOs, CTOs, CAIOs (Chief AI Officers), and their teams across North America and Europe.

Key takeaways from the study are :

[1] AI software spending set to nearly double by 2025: Forrester.

1. Platform Thinking is Taking Hold

Although less than a quarter of the surveyed enterprises have a platform-led approach towards AI, most leaders are actively considering one. Over 70% of the organizations surveyed have already evaluated the benefits of an integrated platform approach. And among these, 84% are exploring or transitioning to platforms. Leaders are moving with urgency, driven by the high cost of new use case development (81% of organizations surveyed) and governance risks of standalone solutions (77% of organizations surveyed).

2. Unified Ecosystems are Seen Unlocking Value

Seven in ten leaders anticipate that unified platforms could reduce costs of new AI use case developments by at least 10%, with one in three anticipating over 15% savings by using Platform AI solutions[2]. Additionally, more than 80% predict faster deployments. These findings underscore the belief of the leaders that moving to a unified platform will directly enhance ROI, innovation, and operational agility.

3. Governance is Now a Design Priority

Over 80% of the leaders report valuing platform-level guardrails to mitigate compliance and regulatory risks. These risks get heightened while managing multiple point solutions. The key concern is not the success of AI, but AI without oversight. This fear of failure of compliance shifts the spotlight on the under-appreciated IT roles such as testing or application management.

4. It is Build and Buy, Not Build or Buy

On the question of building or buying, eight in ten leaders prefer a hybrid approach—buying core platform components while building only for sensitive or proprietary requirements. This balances unique needs such as cost-efficiency, control, and innovation effectively.

5. Success Hinges on Technology and Data Readiness both

Most of the enterprises have initiated the assessment of the platform solutions and are actively transitioning. But only a few organizations—just 1 in 10—claim themselves to be data-ready for AI. Poor data quality is widespread. So, organizations considering a platform-led approach invariably invest to make data accessible, current, and fair. However, data readiness is an endless process, hence knowing when your data is "AI-ready" is equally important.

As enterprise AI scales, the shift towards unified platforms is becoming increasingly clear. And that shift is towards a future where AI systems are not just smarter, but also seamlessly connected, governable, and enterprise ready.

[2] Over and above the savings expected by deploying AI.

Part 01

What does a Unified Platform mean in Agentic AI?

1.1 The Journey from Disjointed Solutions to Unified Platforms

Enterprise automation has evolved from rule-based automation to machine learning, and now to GenAI and Agentic systems (autonomous solutions capable of contextual reasoning and task execution) (See Exhibit 1). Yet, despite their sophistication, most solutions are deployed in silos across departments, leading to duplicated efforts, disjointed user experiences, and rising governance risks. Each point solution, though working with high accuracy now, might pose scalability issues in the future.

EXHIBIT 1

The Evolution of Automation in Enterprise

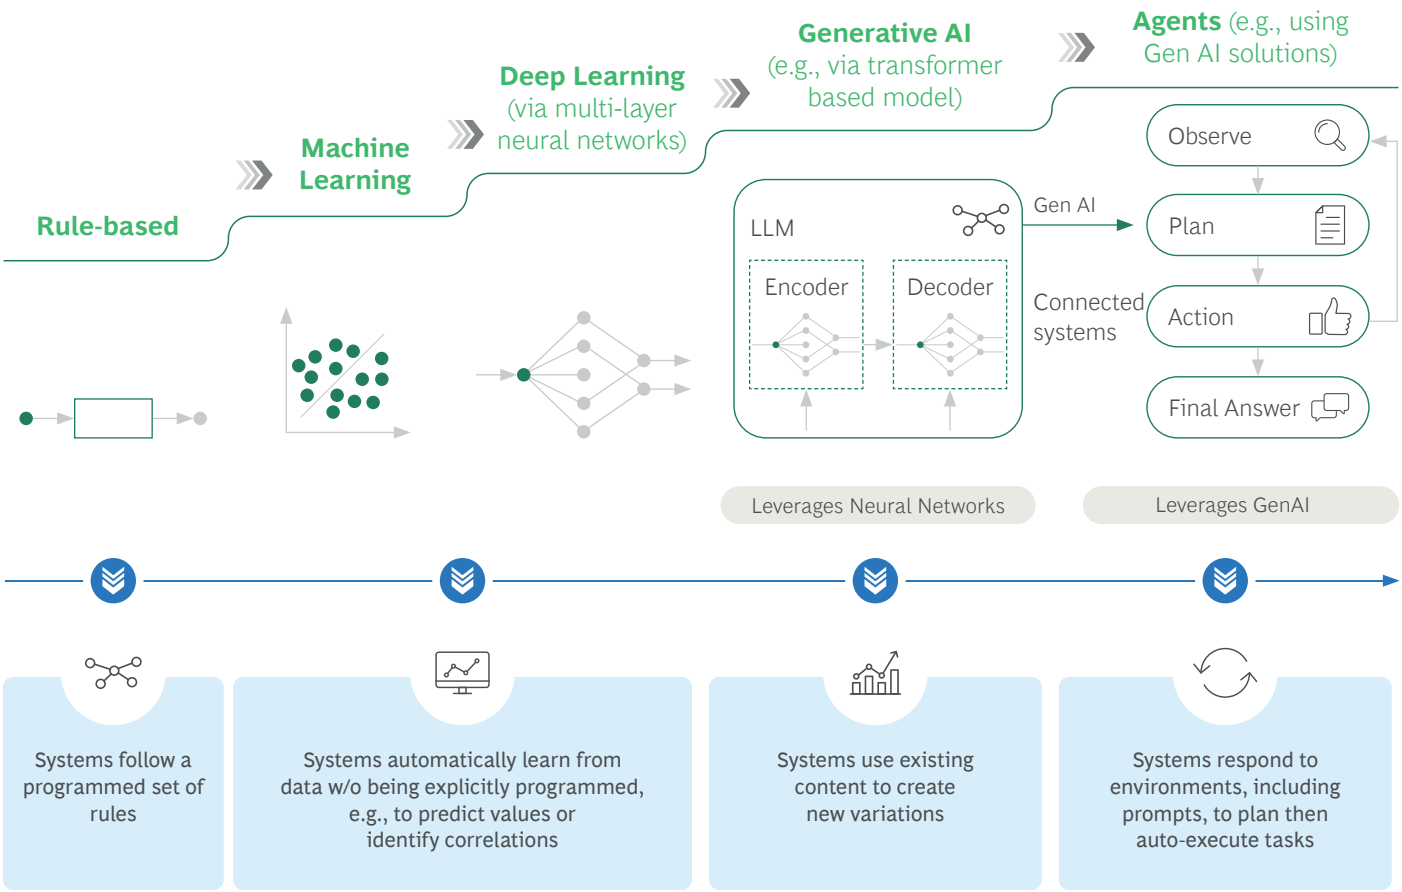

In simple terms, point AI solutions are like cars, designed specifically for commuting, whereas platform AI solutions are like an assembly line, which provide ability to create multiple vehicles for different use cases at scale with standardized rules. Today, fewer than 25% of organizations follow a platform-led AI approach, but over 70% have evaluated a shift to platforms. This transition mirrors past technological shifts. Just as businesses once moved from disconnected spreadsheets to unified CRM platforms for 360° customer insight, or from isolated cybersecurity solutions to integrated systems combining SIEM platforms and precision solutions, AI too is maturing towards cohesion.

While only ~25% organizations follow a Platform-led AI Approach

Q: What approach does your organization follow for deployed AI solutions?

Over 70% organizations have evaluated a shift to AI Platform Solution

Q: Has your organization ever evaluated a platform AI solution approach?

The goal is not to eliminate specialized solutions but to integrate them through a central platform that enables shared governance, data reuse, and cross-functional intelligence.

The path forward is a hybrid model, that coordinates and scales, paired with specialized agents for depth. The destination of smart, interoperable, governable and enterprise-ready AI is finally becoming less hazy.

1.2 The Role of Platforms in A Hybrid Ecosystem

A unified AI platform is an integrated environment where data pipelines, model development, orchestration logic, and governance controls operate within a single architecture (See Exhibit 2). It is built to function as one system, rather than a combination of separate solutions held together by custom workarounds. In the context of Agentic AI, where autonomous agents handle complex tasks and make decisions independently, this type of structure becomes essential. Agents need a consistent foundation to access context, communicate with other systems, and operate across different functions without creating redundancy or friction. The platform ensures these agents work within a coordinated, enterprise-wide framework.

EXHIBIT 2

Advantages of Platform AI Solutions

Illustrative

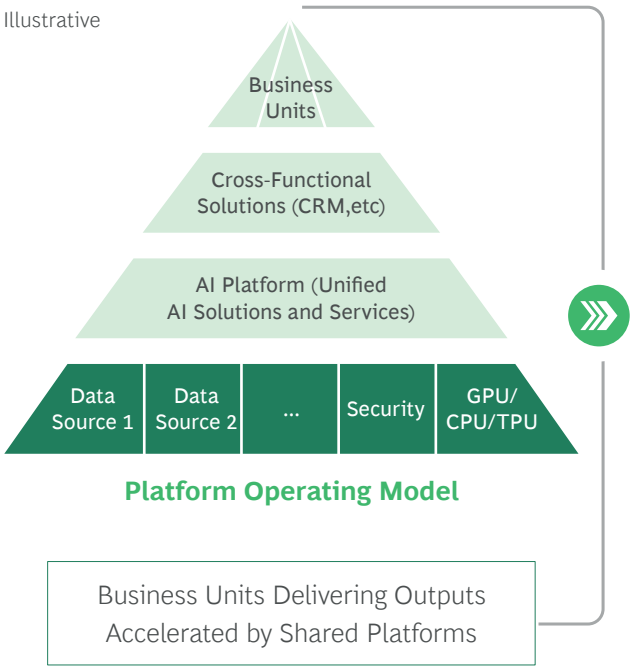

Breaks Silos
By combining capabilities, reducing handoffs and driving E2E outcomes

Removes Duplication
Through the sharing of capabilities across products and BUs

Increases Quality
Through focusing scarce talent and resources, and empowering teams

Accelerates Velocity and Reduces Costs
By providing re-usable, scalable, modular capabilities

1.2.1 Faster Scalability

Reusability of technology and multiple layers of the tech stack may enable organizations to scale faster than non-platform solutions. Leading AI platforms enable assetizing components for reusability across the firm. In the study conducted, over 80% of organizations said they expect faster deployment and greater model reuse (See Exhibit 3).

Nvidia is a key example, where they streamlined their AI initiatives by transitioning from over 65 individual chatbots towards a centralized AI platform. This enabled improved accessibility of AI capabilities of the company at scale[3].

EXHIBIT 3

Expected Scalability from use of Platform AI Solutions

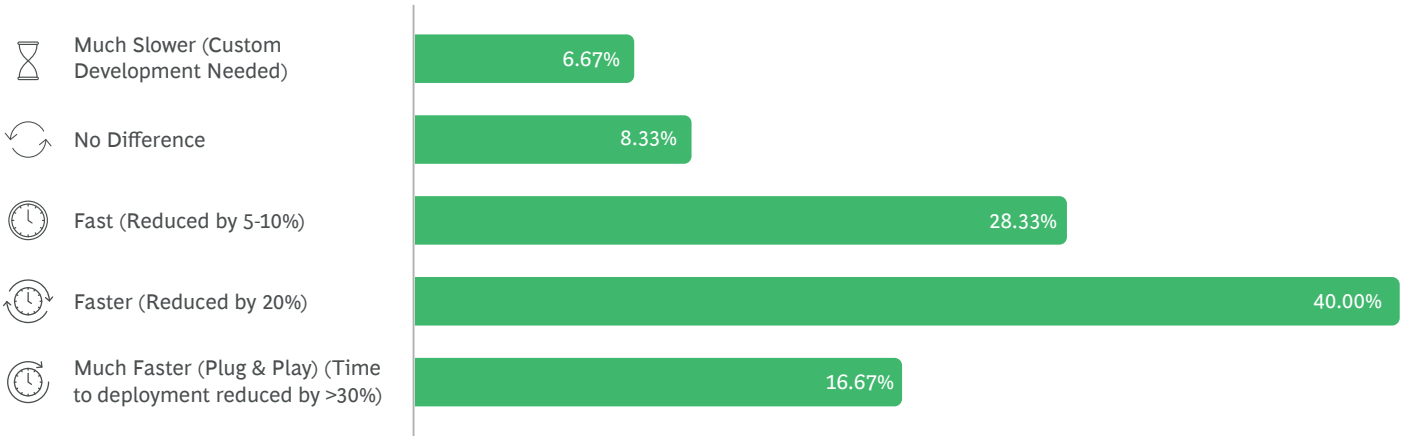

Q: What is the expected improvement in scalability by using Platform AI solution over multiple Point AI solutions?

1.2.2 Centralized Governance and Control

As organizations scale up AI, centralized governance becomes critical to managing risk and ensuring trust. AI platforms enable unified oversight through features like Role-Based Access Control (RBAC), integrated audit logs, and policy enforcement solutions that span all AI projects. The platform acts as an enterprise-wide guardrail, making sure AI development and deployment adhere to the organization's rules and external regulations.

With the advent of platform AI solutions, more than 60% of the respondents estimate that majority of applications and LLMs will be centrally governed via enterprise-level guardrails. This centralized AI governance allows enterprises to skip the lengthy risk evaluations for each individual point solution, reducing decision-making time.

Forrester predicts that by 2025, 40% of highly regulated firms will unify data and AI governance to manage increasing oversight demands. The platform's unified governance "cockpit" helps manage AI's integrity and risks in one place, balancing innovation with oversight.

1.2.3 Cost Reduction and Performance

Over 68% of organizations believe at least a 10% reduction in cost per new use case build by using platform AI solutions[4]. This is majorly attributed to reuse of components and the models for multiple use cases, thereby eliminating the need to build from scratch each time, 25% of respondents believe that all AI models could be reused across multiple use cases while 42% expect significant components such as training datasets, model architecture, and code modules can be reused. A platform pools AI resources, providing "elasticity" to allocate compute where needed. In siloed setups, teams maintain separate (often idle) servers and hit capacity ceilings, whereas a unified platform ensures no hardware sits idle. At the same time, priority workloads always get resources.

Enterprises recognize the value that platforms bring. Out of the 70% of organizations who conducted the detailed evaluation of Platform AI solutions, 84% of these firms are transitioning to platforms underscoring the platform's potential (See Exhibit 4). While this is important and catchy for CIOs, the next step is to develop a clear transition roadmap to align key stakeholders and ensure successful solution implementations. Key concerns from multiple point solutions and potential benefits across segments from Platform AI solutions are mentioned ahead (See Exhibit 5).

[3] HBR: To Scale GenAI, Companies Need to Focus on 3 Factors.
[4] Over and above the savings expected by deploying AI.

EXHIBIT 4

Of the 73% of Organizations that Evaluated Platform Solutions, 84% are Exploring Or Transitioning to Platforms

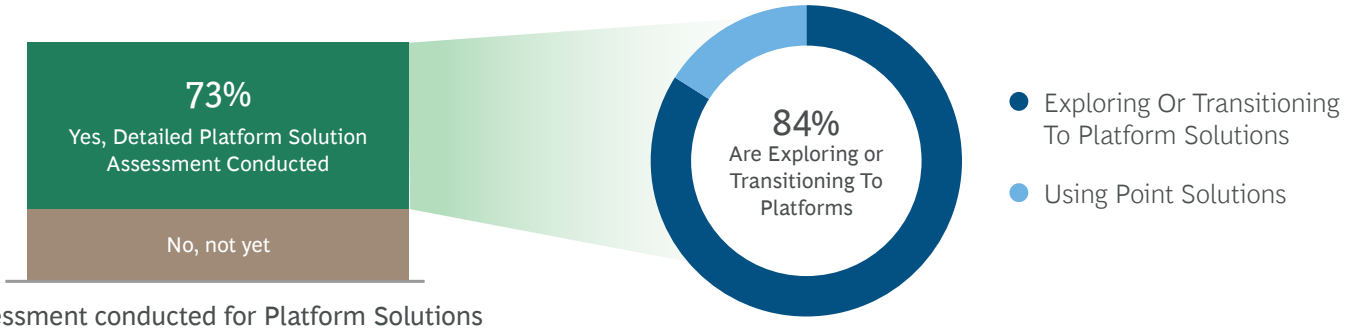

Assessment conducted for Platform Solutions

EXHIBIT 5

Expected Improvement Across Segments By Using Platform AI Solutions

Q: What is the order of importance for your organization in the following parameters with respect to AI solutions deployment?
Q: What are the key hidden costs/barriers to growth that you feel for point solutions?

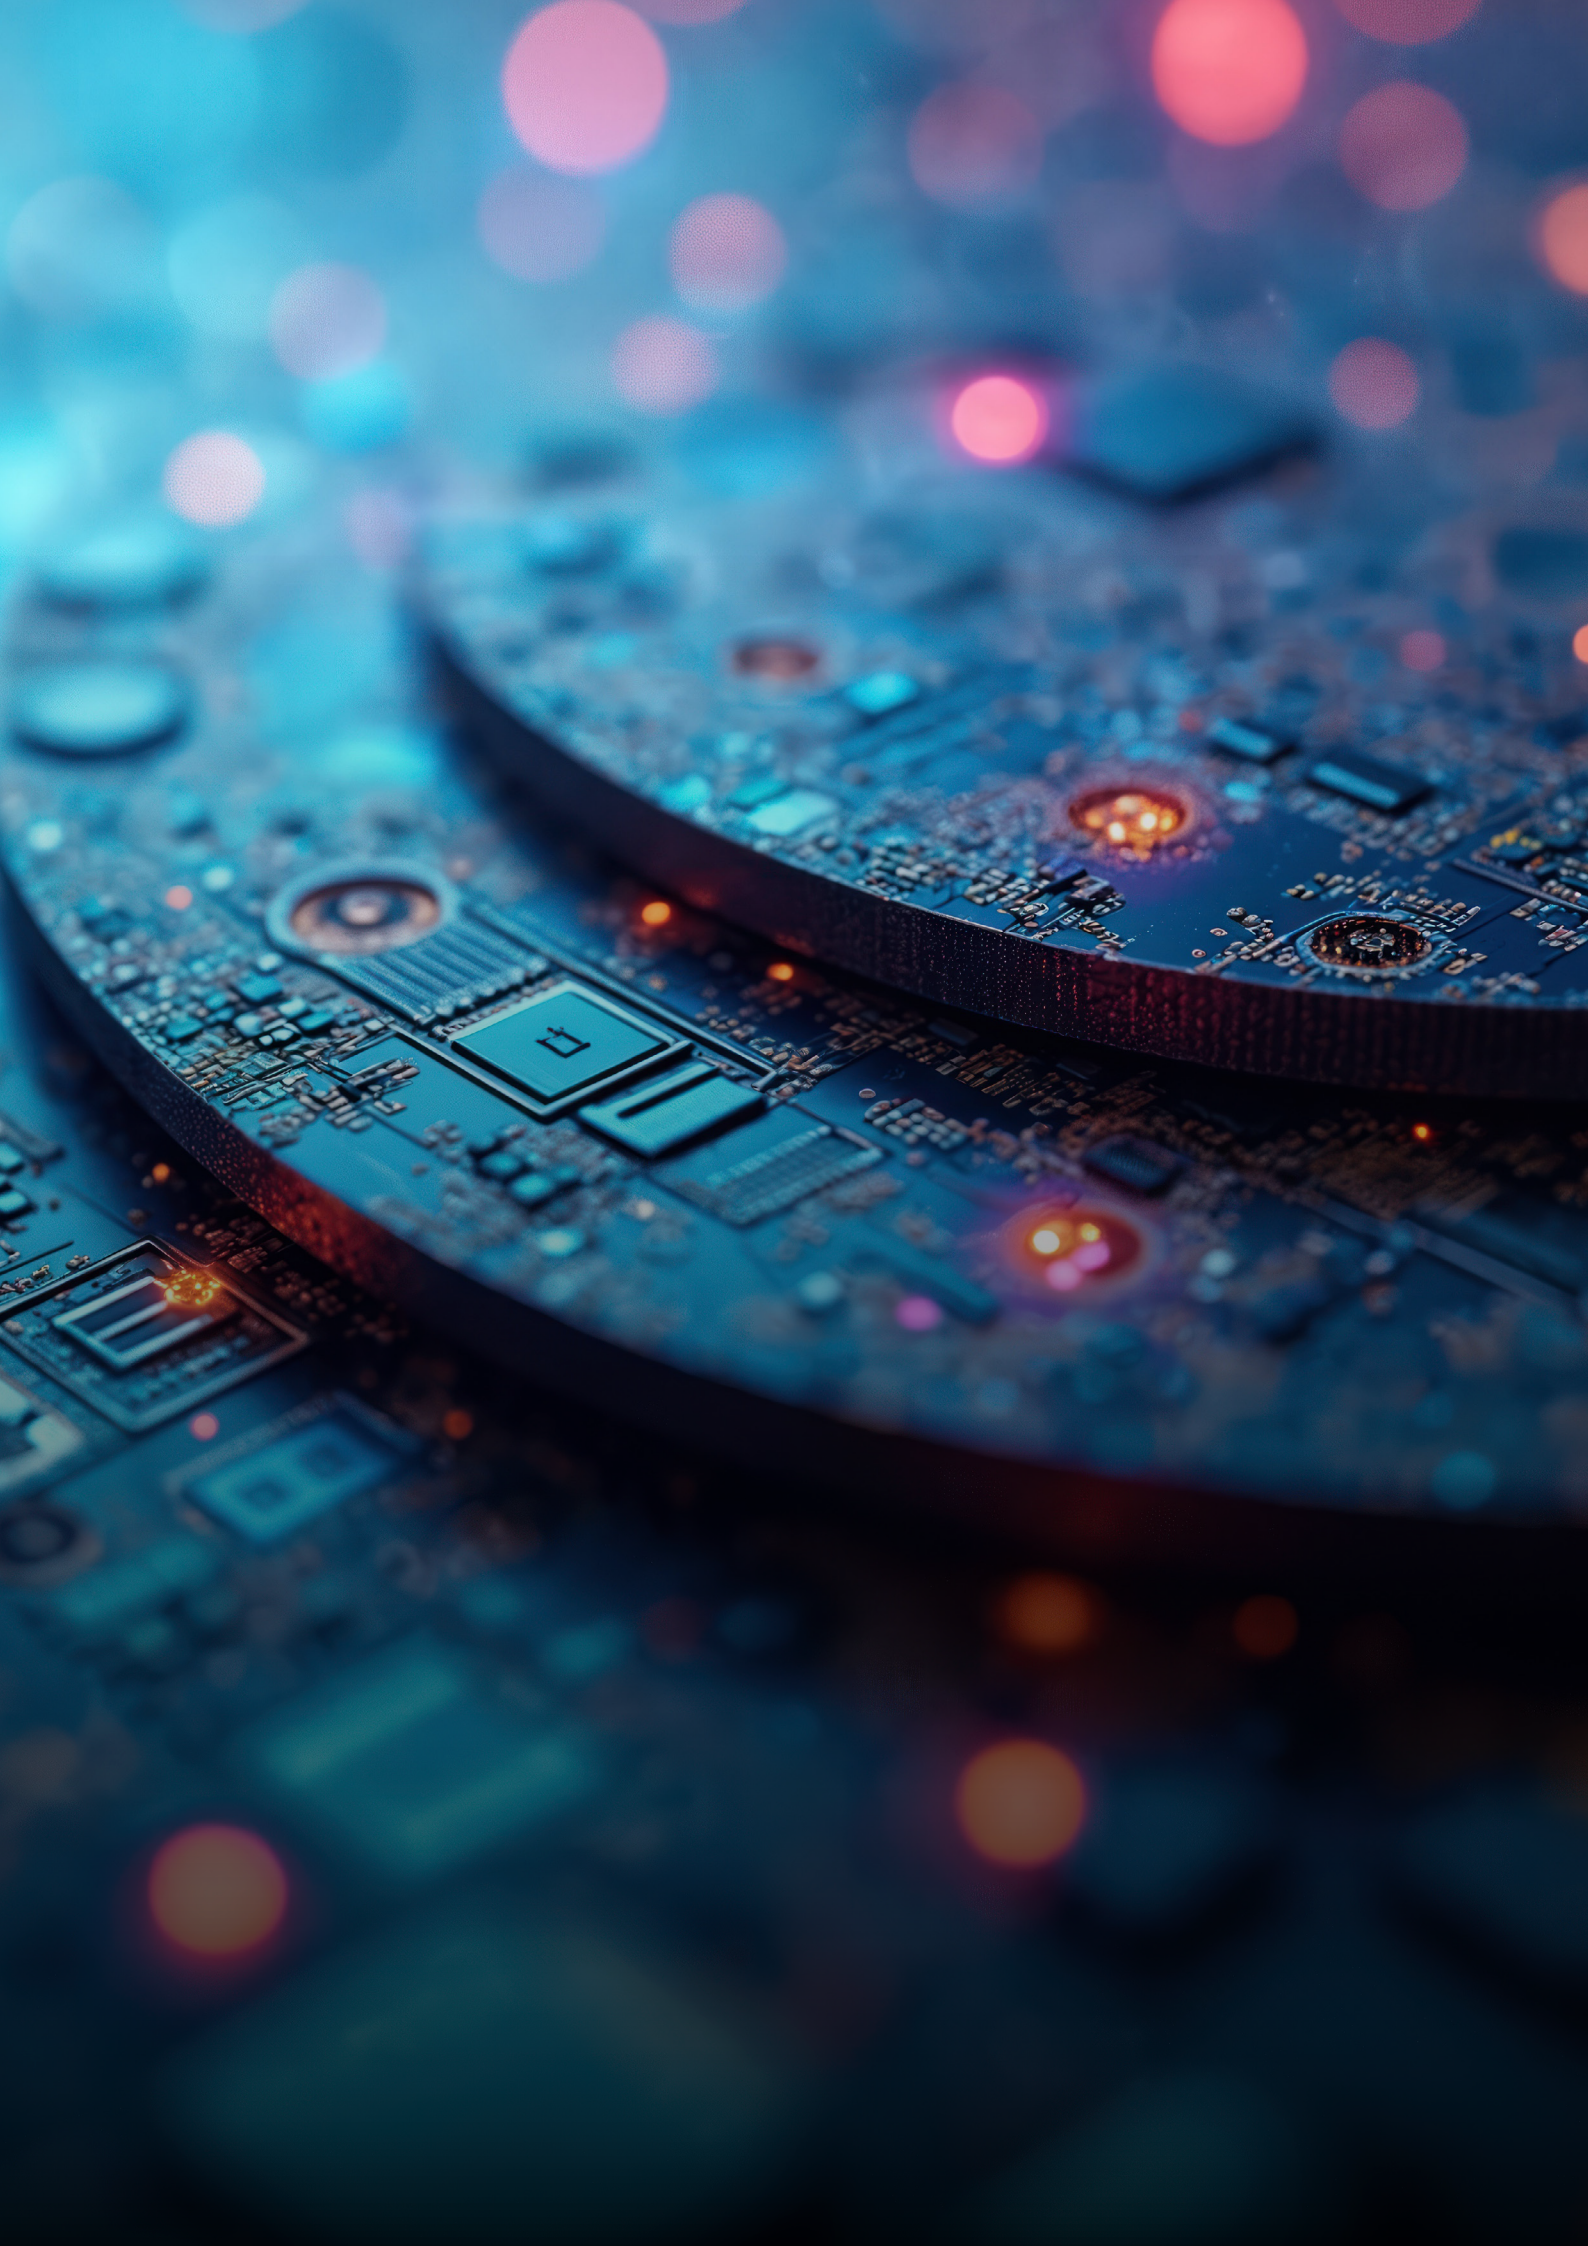

Part 02

Why should you Invest in Building Platform Capability?

2.1. Trade-offs of Disconnected AI

Many organizations take a pragmatic approach while adopting AI across functions. They deploy point solutions to address high-priority issues quickly. The results come at a significant cost. On one hand, there are issues with scattered data governance, and scaling inefficiencies; but on the other, these solutions excel at focused tasks.

EXHIBIT 6

Top Ranked Concerns Faced by Enterprises from using Multiple Point AI Solutions

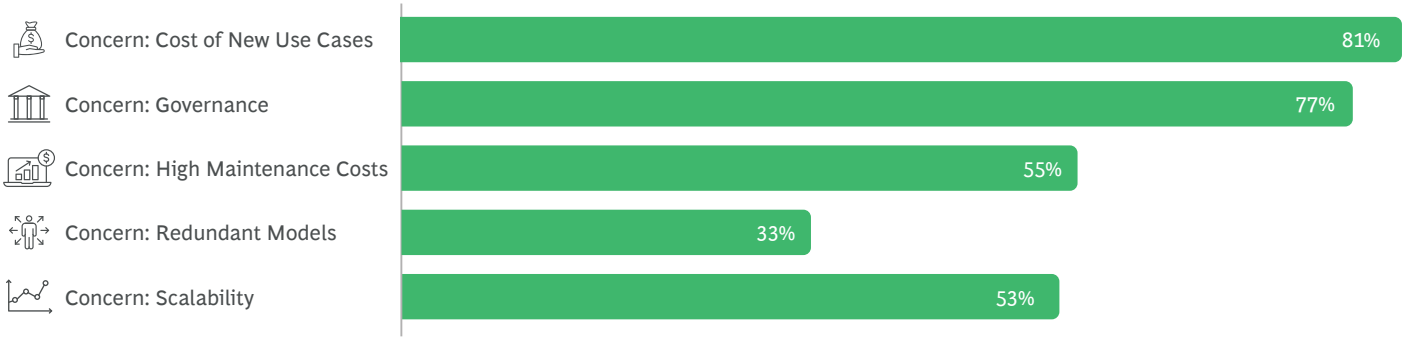

Respondents were asked to pick top 3 concerns

2.1.1 Duplicate Spend: A Hit to ROI

Siloed AI efforts across business units usually leads organizations to solve the same problem multiple times, leading to redundant model development, wasted talent, and overlapping infrastructure. Teams build similar models independently, driving up costs and delaying innovation, while fragmented stacks inflate spending on solutions and services. Thirty percent of AI practitioners cite integration issues from duplicate systems as a top frustration[5]. Disconnected models create inconsistent outputs and require manual reconciliation, which undermines confidence in AI-driven insights.

The cumulative impact combined with implementation costs and a failure to reimagine the operating model might be a significant hit to ROI. Despite surging investment in AI (50% year-on-year), only about 25% of the companies were able to generate value from AI[6]. The ones reaping the benefits are also seeing anemic returns on siloed approaches (as low as 5.9% ROI, below the cost of capital[7]). In fact, organizations that master AI efficiency via centralized, collaborative approaches see up to 3x higher ROI than those taking siloed approaches[7].

2.1.2 A Serious Enterprise Risk

In BCG's latest Global ESG, Compliance and Risk Survey (See Exhibit 7), Gen AI-related compliance was cited as the third most concerning risk after emissions reporting and adverse media coverage. Every additional AI application is a potential point of failure and increases the likelihood of security incidents. Inconsistent or disconnected guardrails create gaps that malicious actors can exploit. Our study uncovers that governance is a top-of-mind concern for organizations using multiple point solutions with over 77% responses (See Exhibit 6). A fragmented approach makes it harder to prove compliance with laws like GDPR[8], HIPAA[9], or emerging AI regulations, as data and decision logs are scattered. Another compliance challenge is the "shadow AI" phenomenon: 88% of organizations reported that employees are using AI (multiple solutions) with or without official approval. These applications, which often

[5] Datarobot: Why AI leaders can't afford the cost of fragmented AI tools.
[6] Latest Build for the Future.
[7] Virtasant: AI Operational Efficiency: Navigating GenAI's True Cost.
[8] General Data Protech Regulation.
[9] Health Insurance Portability and Accountability Act.

are funded by other businesses, are even riskier than their own AI solutions without oversight. As AI accelerates, governance and controls are struggling to keep up.

All these factors emphasize that disconnected AI is not just an IT headache—it is a serious enterprise risk.

EXHIBIT 7

Key Concerns Faced by Enterprises

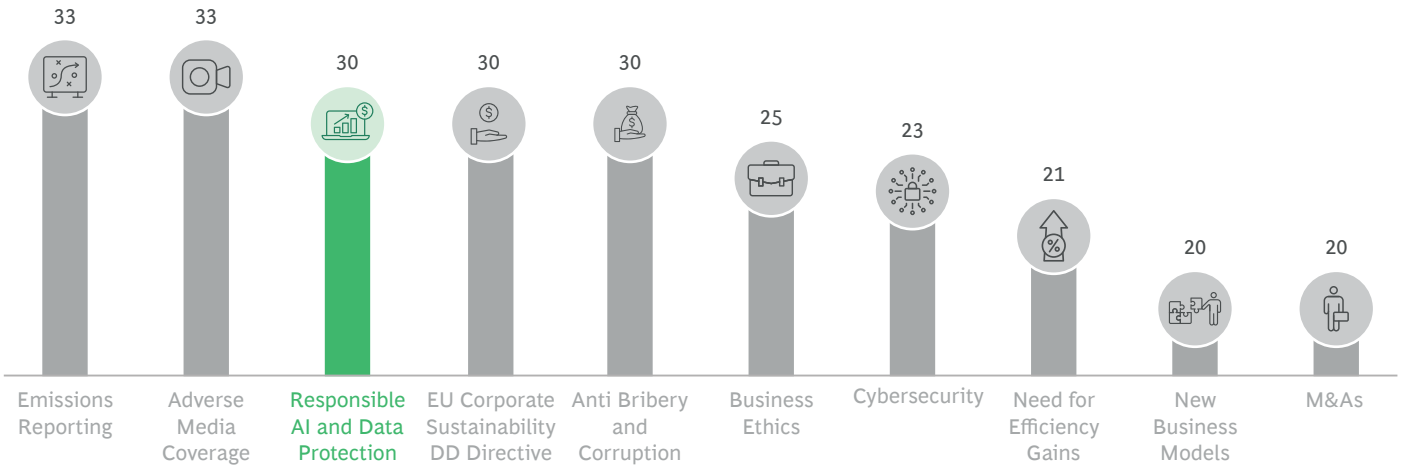

Source: BCG's Global ESG, Compliance, and Risk Survey; BCG Analysis.
Note: Respondents were asked to select their top five risk types.

2.1.3 Duplicate Effort: Maintenance and Scalability Challenges

As much as 15-50% of the cost of GenAI implementation budget may go towards fixing problems associated with difficulty in maintaining and integrating with legacy code[10]. This 50% spend on maintaining operations is taken directly from funding future innovations, hindering future growth.

Over time, the lack of interoperability becomes a barrier to expanding AI capabilities organization-wide. Disconnected AI environments also result in underutilized or redundant compute workloads. Different teams often provide separate servers or cloud instances to run similar models, resulting in redundant compute capacity and operational waste. Without centralized orchestration, it is impossible to optimize performance or infrastructure at scale.

While platform consolidation is gaining momentum, standalone AI solutions designed for narrow, high impact use cases continue to offer real advantages. According to BCG's latest 'IT Buyers Survey', despite a broader shift towards suite-based solutions, CIOs still prefer best-of-breed solutions over platforms when it comes to GenAI adoption primarily for their higher quality datasets, purpose built integrations and fine-tuned models. (See Exhibit 8).

As a result, a hybrid AI strategy is emerging as the preferred model. Leading enterprises are combining the governance, scalability, and reusability of centralized platforms with the agility and precision of specialized solutions. This "both/and" approach enables organizations to balance efficiency with speed—placing platform controls at the core, while retaining the flexibility to deploy targeted solutions where they create the most value.

[10] BCG GenAI Enterprise Foundations Research.

EXHIBIT 8

Best-of-Breed vs Platform (suite) Preference for Multiple IT Solutions

Q: In what categories does your company adopt a best-of-breed vs suite approach?

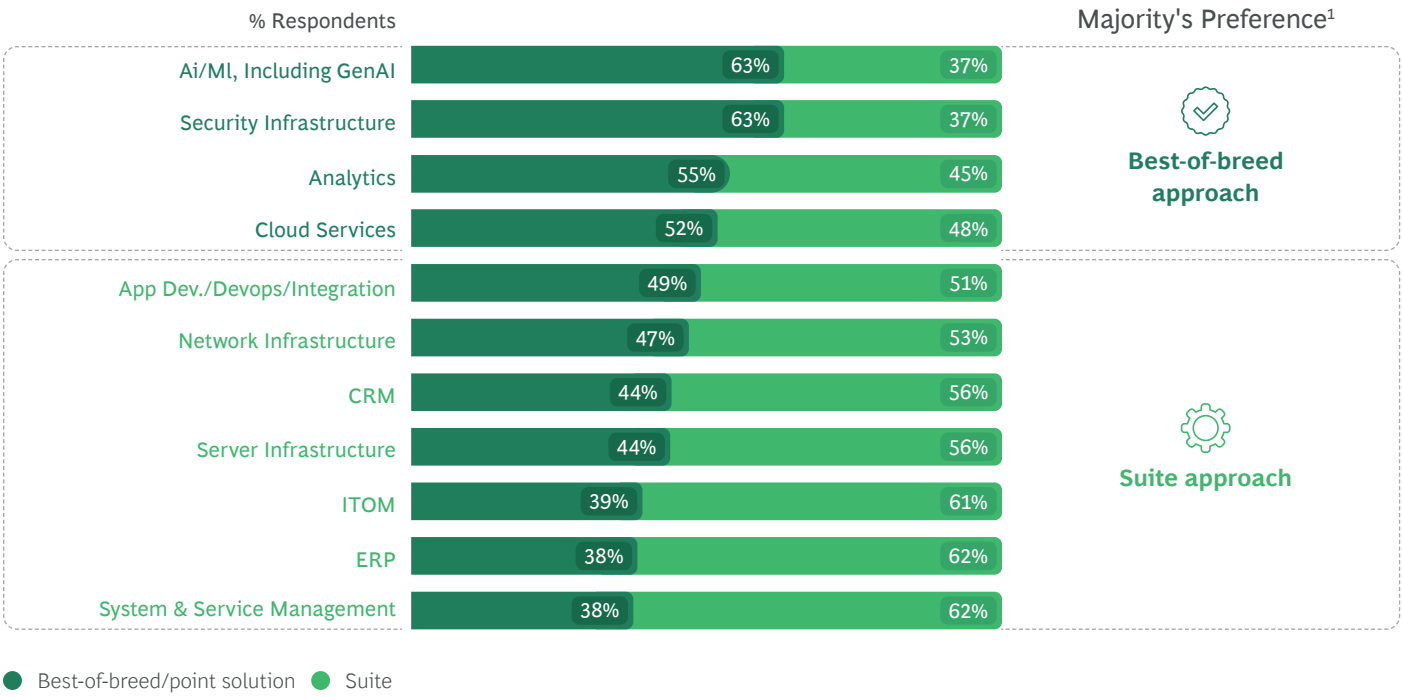

1. Measured on which % of respondents selecting each response is higher than the other.
Source: BCG's Global ESG, Compliance, and Risk Survey; BCG Analysis.

Part 03

How to Build, Scale the Platform and Overcome Roadblocks?

3.1. A Phased Approach to Agentic AI Platform Adoption

Implementing Agentic AI at scale requires more than just deploying individual use cases. A structured platform approach is needed (See Exhibit 9).

The journey begins with aligning infrastructure, defining value, and establishing a clear roadmap. Early investments in architecture, workflow mapping, and use case prioritization lay the groundwork for long-term platform success. Enterprises that approach this phase with a product mindset—designing for reuse, interoperability, and governance—can build a foundation that avoids the fragmentation common in AI deployments.

Following an initial alignment, the focus shifts to controlled pilots that validate business impact and inform enterprise design. Instead of isolated experiments, these lighthouses serve as critical proofs of value, testing platform capabilities and surfacing key integration, data, and adoption considerations. Pilots also accelerate the formation of enterprise guardrails—from risk frameworks to platform governance—ensuring future expansion is both secure and scalable.

Scaling is an iterative process of expanding validated patterns across the enterprise. Successful organizations apply a blueprint-based rollout, continuously improving workflows while embedding platform controls and orchestration.

Change management is an enterprise-wide enabler, requiring leadership sponsorship, employee readiness, and clear articulation of platform value. By taking a phased and platform-led approach, organizations are better positioned to unify AI efforts, reduce cost of ownership, and drive sustained competitive advantage at scale.

EXHIBIT 9

Illustrative Implementation Map for Agentic Platform AI Solutions

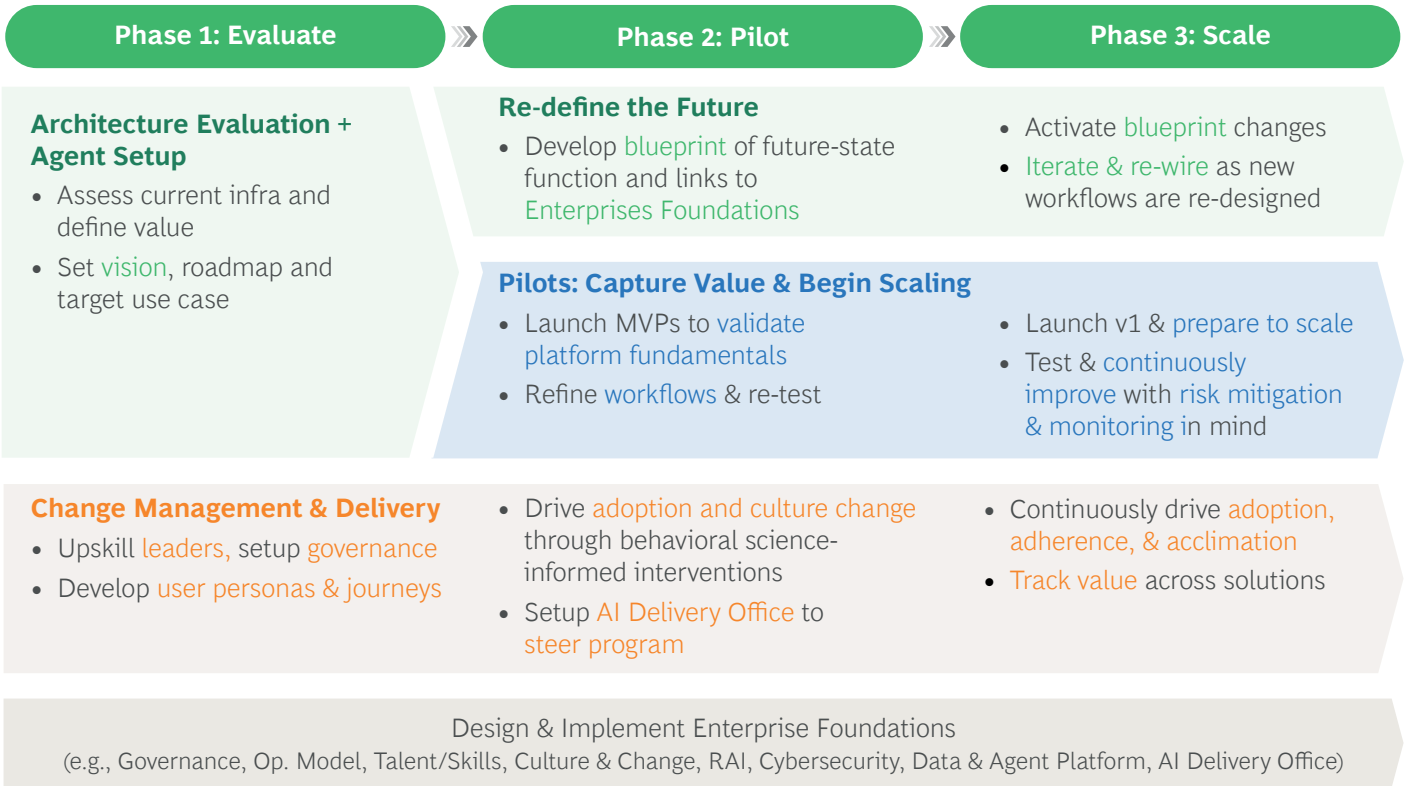

3.2. Build-vs-Buy: Making the Right Decision

The build-vs-buy dilemma is a familiar one for CTOs across evolving technologies. Choosing whether to build or buy Agentic AI platforms is a strategic decision shaped by an organization's priorities, capabilities, and context, and will serve as the source of differentiation needed to stay ahead.

Building is often the preferred path when use cases demand tailored functionality, when data privacy or regulatory sensitivity is high, or when performance thresholds require deep technical optimization. This path also suits enterprises that possess strong internal AI expertise, robust infrastructure, and access to proprietary training data.

On the other hand, buying becomes a more viable route when speed, scalability, and ease of deployment are critical, particularly for well-defined or lower risk use cases. Budget considerations also weigh heavily—while buying minimizes upfront investment, building can offer more cost control over time. Further, if CXOs choose to buy an Agentic AI solution, they need to decide the right provider. This is further dependent on the level of interaction, breadth across the enterprise, the number of systems or tech stack layers involved, and specialization of the task at hand.

In most cases, organizations do not face a binary choice. Instead, a blended approach is emerging as the dominant model, where core, sensitive components are developed in-house, and external solutions are integrated to accelerate time to value. Our study has indicated that nearly eight out of ten respondents have chosen a hybrid approach for setting up platform solutions with in-house solutions used (or leveraged) for data-critical use cases (See Exhibit 10).

This approach, termed as "Buy, Build, and Blend" by Gartner, emphasizes that a mix of in-house and external components often yields the best outcome. This middle path might involve integrating open-source solutions, fine-tuning existing models, or adopting a commercial platform and extending it to enable integration of the specialized point solutions, enabling interoperability and cross-platform orchestration.

EXHIBIT 10

78% of Surveyed Organizations Chose to Build Critical Components of Platforms In-House while Buying Non-Critical Components Off the Shelf

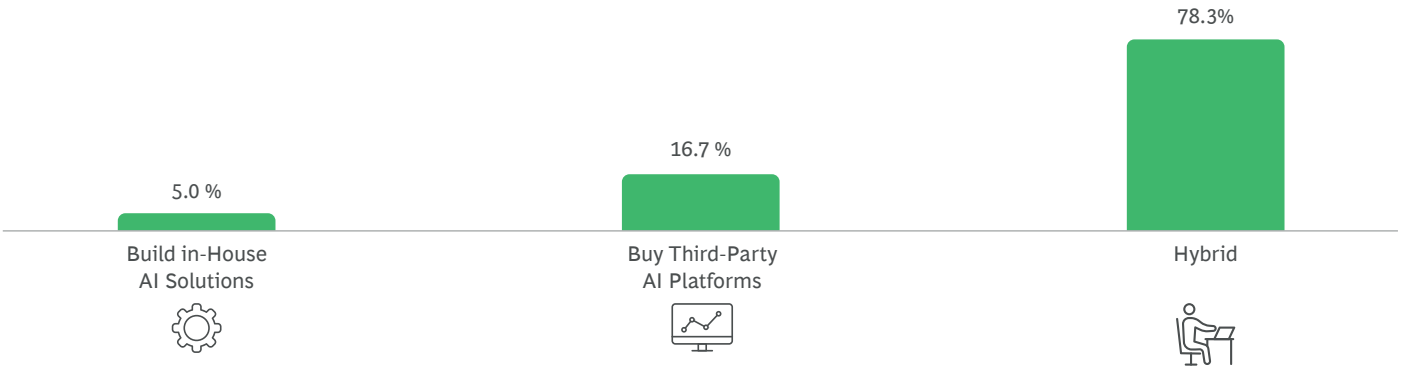

Q: Which best describes your organization's AI strategy?

3.3. Barriers to Platform Adoption—Data Readiness

Enterprise CIOs evaluating AI solutions find that success depends as much on data readiness as on model choice. In fact, data has become the fuel for AI initiatives but only few organizations (1 in 10) consider themselves data ready. Data needs to be high-quality, labeled, fair, accurate and current. Key consideration, however, is also to understand what data readiness means for the enterprises. Data readiness is an endless process, hence enterprises should set up a threshold to define what qualifies as "ready". This decision will be easier when enterprises start with the outcome to be derived from the use cases and gradually iterate and refine.

3.3.1 Fair and Safe Data

Deploying Agentic AI platforms requires that data be both fair and confidentially safe, as these qualities are the cornerstones of an AI system's integrity, scalability, and governance. Fair, unbiased data (with equitable representation across demographics and scenarios) helps eliminate algorithmic bias, ensuring decisions are ethically sound and defensible. Protecting sensitive information (PII and proprietary IP) in AI development is non-negotiable. With such critical data risk, a responsible AI approach integrates human empathy with creativity and work for good. A responsible AI framework across data lifecycle enables clean, fair, and unbiased data for AI (See Exhibit 11).

EXHIBIT 11

Data should be Fair and Safe, Current (Accessible) and Accurate

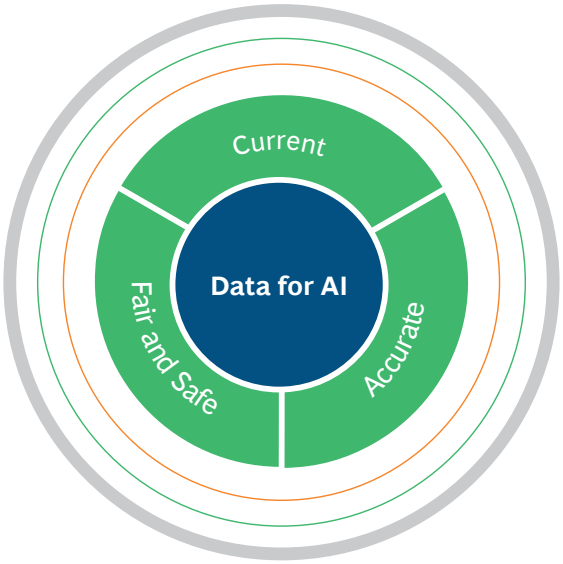

Fair and Safe

Data that is equally distributed without bias and that protects individuals (PII) and sensitive IP

Current

Have the latest and most relevant data

Accurate

Data that represents facts and reality

3.3.2 Breaking Silos via Current Data

Most large organizations store data in myriad systems (CRM, ERP, data warehouses, file shares, and so on), often fragmented by function or geography. These silos inhibit AI from forming comprehensive answers or insights because the model only sees a partial view. While unifying data has not been a new concept, advent of cloud, visualization technologies, reduced storage and compute cost and standard APIs have made the exercise easier.

Industry experts predict a shift away from application-specific systems of record towards unified data hubs that provide a single, contextualized view of enterprise data for AI purposes. CIOs should assess where unified data will yield a clear upside.

Many firms start by unifying high-value domains, consolidating and pulling relevant information (for instance, integrating sales and customer support data for service AI) while expanding capabilities as their data governance matures. For instance, over 98% of Morgan Stanley's advisor teams now use an internal GPT-4-powered assistant for "seamless internal information retrieval," dramatically reducing the time to find answers in silos of PDFs[11]. A pragmatic interim approach is to use virtual unification—that is indexing disparate sources in a vector database for semantic search, without physically moving all data. This approach, a form of RAG architecture, brings together information at query time.

3.3.3 Need for Accurate, High-Quality Data

In supervised learning and fine-tuning scenarios, curated datasets of inputs and expected outputs are indispensable for teaching a model domain-specific knowledge. Studies show that quality, diversity, and relevance of data are critical. High-quality data can improve factual accuracy and reduce biased or "hallucinatory" outputs by anchoring the model.

While data readiness is of critical importance, CIOs need to be cognizant of when data is truly "AI- ready" based on business requirements and use cases. They should start with the quality and quantity of data needed. Organizations may not need extensive labeled data if their AI strategy leans on existing pre-trained models or unsupervised techniques. Using a large public LLM via RAG might help bypass the need for a huge, labeled training set – instead, the model can be kept generic and fed with relevant enterprise content at query time. In such cases, data quality still matters (the content fed to the model must be accurate and vetted), but it does not require the upfront labeling effort of fine-tuning.

However, if an organization pursues a custom model or extensive fine-tuning, investing in labeled data becomes non-negotiable. The need also varies by organizational readiness: a mature data organization with robust DataOps pipelines and annotation workflows can more easily supply high-quality data, whereas a less data-mature firm might pilot GenAI on smaller, existing datasets first.

Our study shows that data readiness is no longer an impediment to AI investments but a parallel focus (See Exhibit 12).

[11]OpenAI: Morgan Stanley uses AI evals to shape the future of financial services.

EXHIBIT 12

Organizations have Evaluated Platform AI Solutions Regardless of their Data Readiness Level

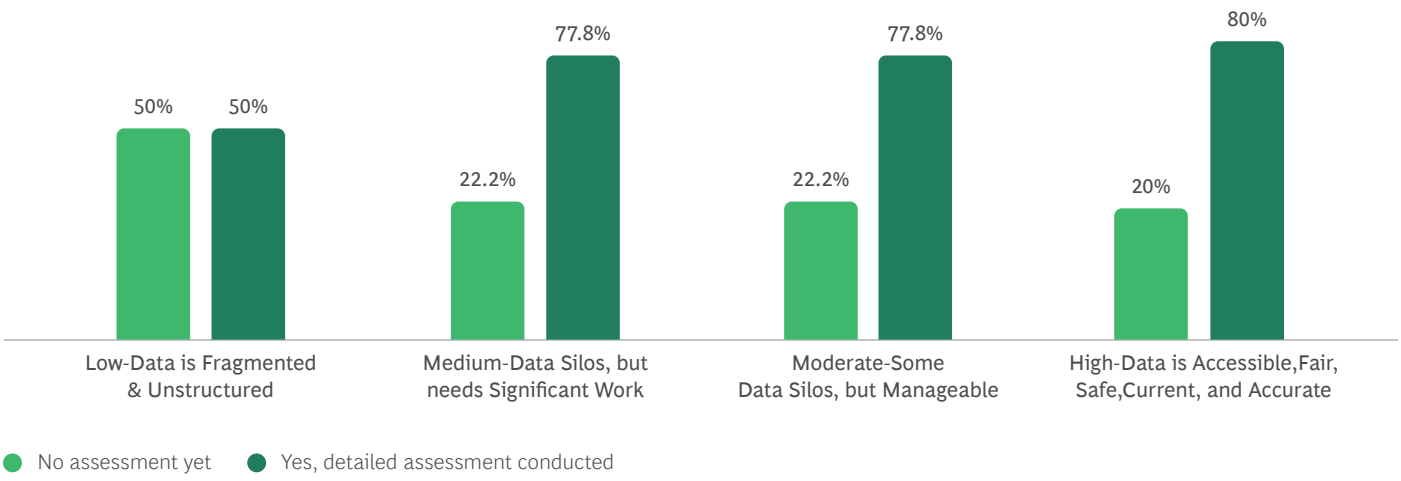

3.4. The Path Forward: Unified Platform, though still a Distant Reality

As enterprise AI scales, the shift towards unified platforms combined with specialized AI solutions appears to be moving. This shift is towards a future where AI systems are not just smarter, but also seamlessly connected, governable, and enterprise-ready. Fragmented deployments are difficult to scale and difficult to govern but provide accurate outputs for specific use cases.

Striking the right balance between specialized, nimble point solutions and comprehensive, scalable platforms allows companies to achieve precision and agility without sacrificing governance and operational efficiency. A hybrid platform approach enables interoperability, accelerates innovation, and embeds trust through consistent governance and oversight. But this path requires intentional choices: what to build, what to buy, and how to ensure data is ready to power AI at scale.

Success will come to those who move deliberately from fragmented experiments to enterprise-grade solutions where point solutions and platform solutions coexist. The destination is shared: connected, governable, and scalable AI—ready to drive the next wave of business transformation.

About the Authors

Authors

Vikash Jain is a Managing Director and Senior Partner at Boston Consulting Group, and a leader in BCG's Tech Practice. He leads the tech services practice for BCG

Shaleen Sinha is a Managing Director at Boston Consulting Group and heads their Growth Tech Practice in India which works with late stage digital natives. He is an ex-CEO of Zivame

Amit Kumar is a Managing Director and Partner at Boston Consulting Group, focused on Tech Products, Services and helping organizations scale up their GenAI offerings commercially

Nicholas Clark is a Partner and Associate Director at Boston Consulting Group, and brings hands-on experience of implementing AI tools across enterprises

Prasanna Arikala serves as the CTO and Head of Product at Kore.ai, driving product strategy, technology vision and innovation. He is a founding team member at Kore.ai

Cathal McCarthy is the Chief Strategy Officer at Kore.ai, and has deep expertise in CRO, COO, CCO, and CMO roles within Fortune 100 companies and high-growth startups

Cobus Greyling is key evangelist at Kore.ai, and a specialist in Language Models, AI Agents, Agentic Applications, Development Frameworks, and Data-Centric Productivity Tools

Benjamin Quirt is Vice President at BCG X, Boston Consulting Group's tech build and design unit, with experience in large scale sales and service transformations

Key Contributors

Saurav Mohanty is a Principal at Boston Consulting Group.He specializes in growth acceleration and offering development, helping technology companies scale effectively and innovate in an increasingly digital world

Shreya Anil is a Project Leader at Boston Consulting Group with over 10+ years of experience across digital transformation, go-to-market, product strategy and risk management for technology companies

Anoop Gopal Singh is a Consultant at Boston Consulting Group, with experience across Technology Media and Telecommunication, and Growth Tech space at BCG

For Further Contact

If you would like to discuss this report, please contact the authors.

Acknowledgments

We extend our gratitude to the BCG Global and India Marketing Team for managing the marketing. We would also like to thank Saroj Singh, Abbasali Asamdi, Ratna Soni, Soumya Garg, and Kshama Sawant for their contributions to the editing, design, and production of this report.

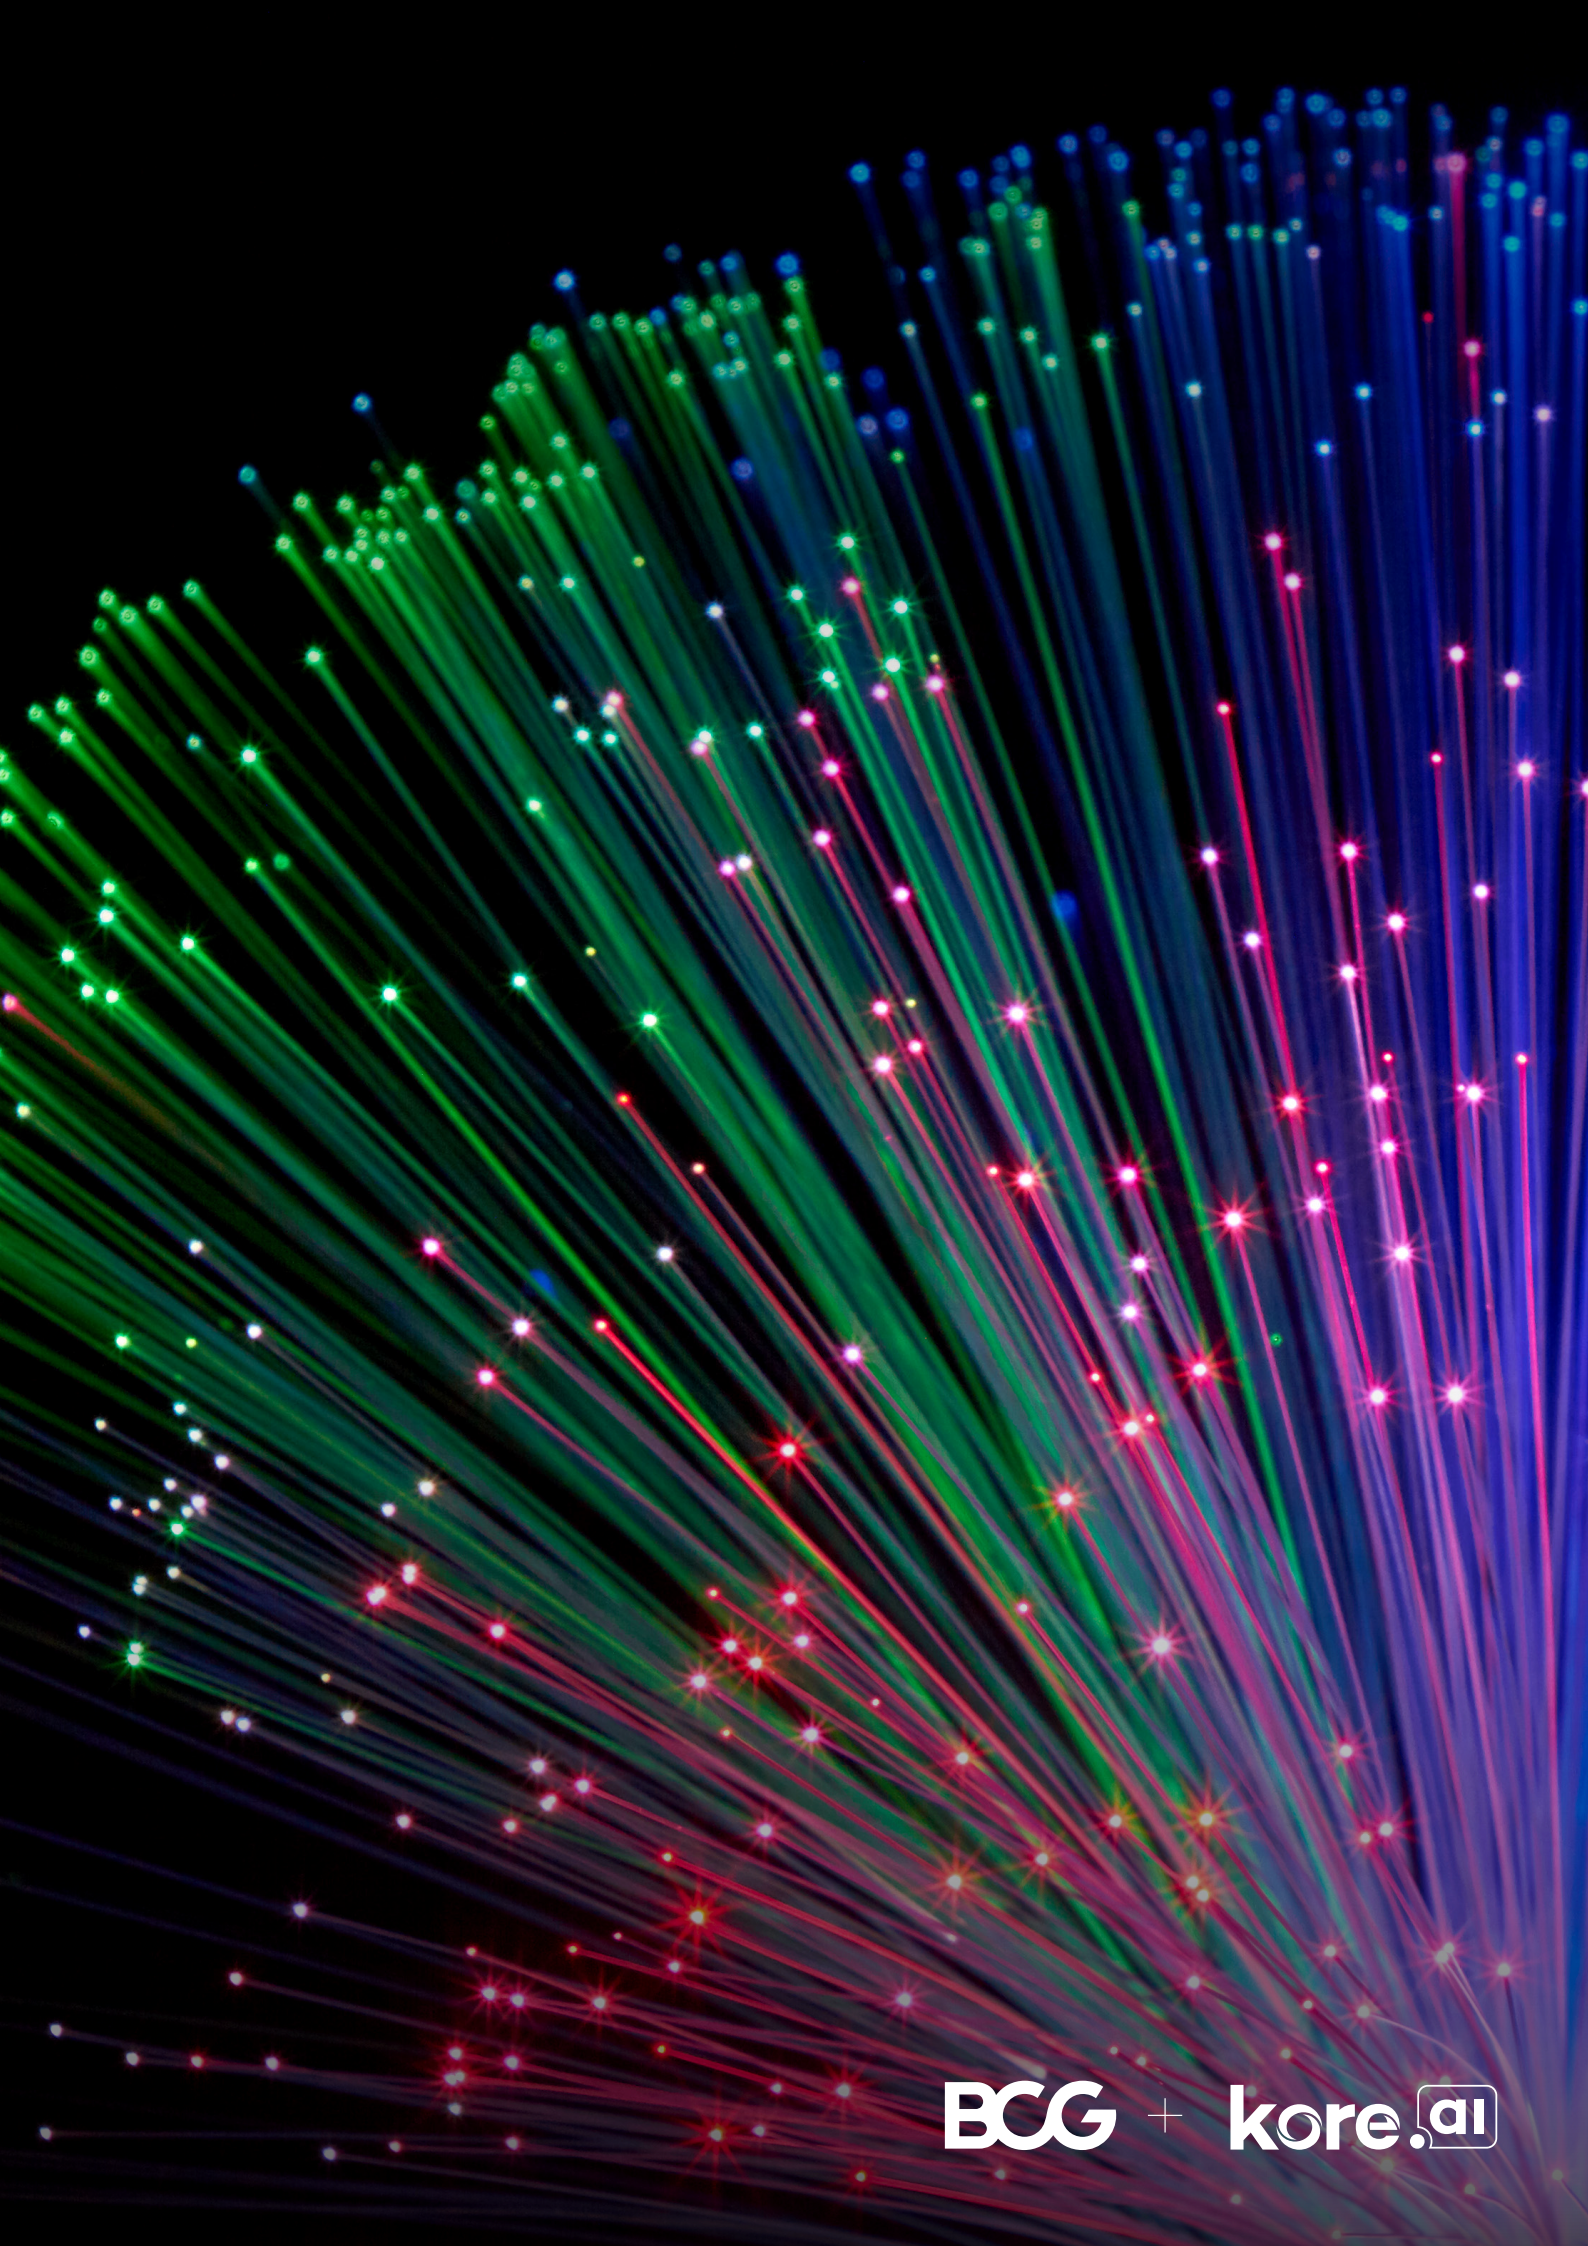
BCG + kore.ai