



ARTIFICIAL INTELLIGENCE

How to Create AI Value by Balancing Innovation and Risk

By [Matthew Kropp](#), [Julie Bedard](#), [Ulrike Schwarz-Runer](#), [Steven Mills](#), [Anne Kleppe](#), [Lisa Kray](#), [Marc Roman Franke](#), and [Zeeshan Shah](#)

ARTICLE OCTOBER 26, 2026

High-profile security incidents involving frontier AI models are keeping many executives up at night as they show agentic AI's unpredictable nature. These events have largely stemmed from training runs of not-yet-released models, but AI is increasingly giving rise to sometimes costly hallucinations, as well as bias and privacy violations, via public-facing chatbots. Leaders are also mindful that employees can introduce risk when they cognitively offload decision making to AI for business-critical tasks. All told, this environment is sparking a global debate among senior


leaders: Can they encourage employees to experiment with AI and scale impactful innovation while managing the risks?

On the basis of our latest research, we believe they can. BCG's 2026 Applied AI Index survey canvassed 1,330 leaders of global enterprises and found that companies that scale AI innovation alone see twice the ROI compared with those that have low levels of AI innovation. Additionally, companies that scale innovation alongside AI governance see three times the value. (See Exhibit 1.)

EXHIBIT 1

ROI Scales with Innovation and Accelerates with Governance

AI addressable value



Source: BCG's 2026 Applied AI Index.

Note: n = 1,330. Value is the addressable revenue plus cost savings calculated thus: the sum of realized revenue growth times the addressable revenue baseline plus cost savings times the addressable cost baseline from AI. The trend shown was calculated from a regression model and cluster analysis and pressure tested with univariate and multivariate effect comparisons. Thirteen percent of companies scaled governance and accountability without scaling innovation, achieving an ROI of 1.3%. We determined companies that scale innovation from a composite measure of the following: use of frontier tech scouting, AI use in research and ideation, AI use in product development, workflow reimagining for AI, development of products that integrate AI, and strategic clarity of AI ambition. We determined companies that scale governance from a composite measure of whether they have implemented the following: responsible AI, data governance, data compliance and protection, zero-trust cybersecurity, and agentic governance controls.

Yet, this value isn't being widely captured. Executives tell us that they want to use AI to create a culture of innovation and that they recognize the next big idea could come from anywhere in the firm. They also understand that the person who knows how to optimize a workflow is someone who is executing that workflow every day. But in a heightened risk environment, leaders also realize that a surge in scaled AI adoption and citizen development isn't straightforward. (See "What Can Go Wrong? An Overview of Risks When Everyone Uses Agentic AI.")

– What Can Go Wrong? An Overview of Risks When Everyone Uses Agentic AI

AI does not create new categories of enterprise risk. Cybersecurity, data and privacy, technology and operations, people and workforce, conduct and compliance—it's the same list that a risk function already manages. What changes is the radius, amplified by three properties of agentic systems.

- **Breadth.** At a system level, one defect spreads across every workflow it touches, which means that a single customer-facing error has the potential to affect thousands of customers rather than one. In addition, the use of shared models, plug-ins, and vendors means that a failure in one place cascades into others, leading to failure chains and difficulty in tracing the root cause. At an individual level, anyone with a general-purpose AI model can build tools, offload important decisions, and automate a business-critical workload that requires a human in the loop. Tools can get deployed outside centralized IT channels, bypassing a security review and possibly even gaining access to sensitive data.
- **Speed.** Agents act, and erroneous outputs spread and scale errors, bias, and other issues faster than humans can detect or catch them.
- **Opacity.** As systems become more autonomous, their decisions get harder to explain—and it becomes unclear who, or which system, acted. Furthermore, the same input can produce different outputs. Testing alone cannot prove reliability, and nondeterministic behavior evades signature-based detection.

The consequence is a steepening incident curve. The AI Incident Database documented that reported AI incidents reached 445 in 2025 from 106 in 2022, and 186 have already been logged in 2026 year to date. Moreover, AI risk does not require a bad actor. In fact, few recent public incidents had an intentionally bad actor. Most resulted from an agent doing what it was asked to do but using a path nobody anticipated and with nobody watching. That is why rigorously testing the tools and designing technical controls for the build phase are insufficient: many reported incidents involve a system behaving correctly by its own logic, in production, after review had already been passed.

It is very difficult to identify the most impactful agentic workflows developed by individual employees or to scale agents' development and adoption to capture meaningful value. Moreover, when we surveyed an additional 142 leaders who develop AI policy in an organization with scaled AI deployments, they report that, on average, 16% of their employees use unapproved tools because the sanctioned tools are too restrictive or they are unaware of their options. The key question for executives, therefore, isn't *if* they can strike the right balance between innovation and risk, but *how*.

Start with Transparency and Accountability

To better understand how enterprises manage AI risk, we combined our analysis of survey data with interviews of more than 30 executives from organizations that are at the forefront of scaling innovation—including leaders in our own firm. (See “Our Methodology.”) Overall, our data identified two critical levers for scaling governance alongside innovation.

– Our Methodology

In this article, we drew from three primary sources of data: a survey conducted by the BCG Institute, the 2026 Applied AI Index survey conducted by BCG, and a series of interviews conducted by BCG with experts in AI accountability at BCG and other organizations.

BCG Institute surveyed 142 senior leaders at large companies from September 2, 2026, through September 11, 2026. Every respondent works at an organization that has more than \$1 billion in annual revenue and that has deployed or actively piloted AI in at least one business process. Most respondents have established AI programs: 52% run more than 20 AI use cases in production and 26% run more than 50.

Every respondent is close to AI governance for their organization: 25% set AI policy or make the final decisions on it, 66% help shape it or sit on the body that approves it, and 8% know in detail how it is applied in practice. Respondents are directors (42%), vice presidents (27%), senior vice presidents (18%), and board or C-suite members (13%). More than half are in technology or IT or data and analytics functions; the rest work in strategy, sales and marketing, operations, software, risk and compliance, finance, HR, and legal.

Companies are headquartered in the US (64%) or EU (36%), and 85% operate globally. Eight sectors are represented: technology, banking, and consumer goods and retail (25 each), health care and life sciences (21), insurance (19), fintech (12), biopharma (11), and telecommunications (4). In addition, 33 companies are

cautious adopters, 25 are accountable innovators, 59 are unsupervised adopters, and 17 are transparent experimenters.

The 2026 Applied AI Index survey canvassed 1,330 firms in a global demographic in the following sectors: consumer; energy; financial institutions; health care; industrial goods; insurance; public sector; travel, cities, and infrastructure; and technology, media, and telecommunications.

Finally, we relied on 30 interviews with experts and senior leaders involved in AI outcomes at organizations for the examples throughout. These leaders are from large businesses (more than \$1 billion in revenue) that are based in the US or EU and that represent a breadth of industries, including banking, retail, insurance, biopharma, and technology.

The first lever is transparency. Leaders need to know how AI is being used and what agents are doing in their organization, ideally in real time. To determine organizations' level of transparency, we assessed if respondents in our survey of 142 leaders have three control mechanisms: an inventory of AI systems and agents with a named product owner for each, systemic visibility into the actions of AI agents and the people using them, and the capability to pause or shut down AI systems quickly if something goes wrong. Of course, these mechanisms don't represent systemic transparency around citizen development, which is a significant and unsolved challenge, but they do enable us to measure the adoption of emerging capabilities for enterprise-level transparency.


The second lever is accountability. Every decision on AI tool design, deployment, and usage needs to have a clear owner, and this accountability over decision making must start with executives. On the basis of our survey questions, we assessed if organizations report having central executive accountability for AI decision making, formal documentation of accountability, a single or core responsibility framework, and accountability that expands to the end user. We also evaluated whether respondents know specifically what they are accountable for.

If we look at the two levers in a matrix, we see four company archetypes emerge. (See Exhibit 2.)

EXHIBIT 2

Governance and ROI Require Transparency with Accountability

Median AI value impact by quadrant



Source: BCG Institute's 2026 AI Governance Survey.
Note: n = 134; 8 responses were removed from the analysis because they did not answer the survey questions on ROI. Value is an increase in revenue as a percentage of total annual revenue plus a decrease in annual costs from AI as a percentage of total cost base. Survey questions about accountability included ones about how decisions are made and documented, who is accountable, and how well that accountability is documented. Survey questions about transparency included ones about visibility into what agents exist and the actions of the agents, while maintaining control.


With low accountability and transparency, *unsupervised adopters* have little central control over innovation. The enterprises in this quadrant are the least likely to document decisions and track questions or requests that get escalated to senior levels. They are also the least likely to have clear protocols for preventing employees from building unapproved tools, formal AI-risk requirements for procuring third-party tools, or a single place to send all use-case questions.

Transparent experimenters are almost 50% more likely than companies in any other quadrant to give general-purpose AI models—such as ChatGPT or Claude—to all employees by default. When compared with *accountable innovators*, transparent experimenters are more than three times as likely to default to a proceed-unless-told-to-stop approach when AI use case approvals come up against a project deadline. At the same time, nearly 50% do not require AI training for employees and 65% do not extend accountability to end users. This may, in part, explain why 24% of the respondents in this category report that more than 20% of employees use unapproved tools (so-called shadow AI).

Cautious adopters are half as likely to have 50 or more AI use cases in production than accountable innovators have. (See Exhibit 3.) Employees of cautious adopters are often afraid to innovate. Alternatively, they innovate by neglecting their accountability, with 46% of these organizations estimating that 20% or more of their employees are using unauthorized tools. While cautious adopters' speed of approvals is faster (55% approve AI use cases in under four weeks versus 44% for the accountable innovators), they may be wasting time on unnecessary approvals: 21% claim that more than half the escalations to a central function in the past quarter could have been answered with existing guidance.

EXHIBIT 3

Accountable Innovators Scale AI Use Cases with Appropriate Oversight



Source: BCG Institute's 2026 AI Governance Survey.
Note: n = 140; 2 responses were removed from the analysis because they did not answer the necessary questions.


Six Steps to Accelerate Innovation Alongside Governance

It is no small feat to foster a culture of innovation that is anchored in accountability along with governance and processes that optimize for transparency. (See Exhibit 4.) By researching the accountable innovators, combined with our collective experience helping leaders achieve these goals, we've identified six practices that make accountable innovators stand out from the crowd.

EXHIBIT 4

Accountable Innovators Lead Across Culture, Accountability, and Governance

RESPONDENTS WITH EACH PRACTICE IN PLACE (%)



Source: BCG Institute's 2026 AI Governance Survey.
Note: n = 142. Each row corresponds to a specific question asked in the survey and displays the response frequency by quadrant.

Fostering a Culture of Innovation, Anchored on a Zero-Trust AI-Engagement Model. Many leaders are still struggling to encourage the impactful adoption of agentic workflows, let alone the responsible use of AI. Simple things, such as overly anthropomorphizing AI, can cause users to rely too much on AI and reduce their willingness to take accountability for their work.

In our interviews, we found that accountable AI leaders encourage everyone to build and use AI in a way that is reminiscent of the cybersecurity concept of zero trust. Zero-trust cybersecurity operates on the principle that no user or system should be automatically trusted. Instead, continuous authentication, authorization, and validation of security configurations are required before access is granted to applications and data.

This translates well into the AI space. For example, every user should question whether it is appropriate to use AI for a task or decision given that context can change the risk level. Every output from general-purpose AI frontier models (such as Claude or ChatGPT) should be verified with the assumption that something is wrong, biased, or worse. And every citizen-built agentic workflow in a business-critical process should be thoroughly vetted before it is used. Taking such steps helps to prevent the propagation of bias, the spread of vulnerable code, and the illegal use of sensitive data. They also help ensure decisions are not being improperly passed on to AI.

Building the Necessary Competencies Across the Organization. Zero-trust engagement with AI is not easy and requires new competencies across the board. Role-specific training is a table-stakes best practice that is also required by the EU Artificial Intelligence Act. Three-quarters of our

respondents offer general training to all employees, and 70% provide training on responsible or ethical AI use.

What distinguishes accountable innovators is the design of learning paths that infuse their desired cultural attributes and accountability models into their employees' daily work. Accountable innovators set the expectation that AI should be used to improve performance, but they do not make AI models and tools available to all by default. Instead, these leaders are one-and-a-half to two times more likely to require employees to complete training on AI tools, take responsible AI training, and have role-specific training. These businesses are also more likely to define a standard of what employees are expected to know about AI through an AI fluency ladder or a competency framework.

Keep in mind that responsible AI only starts with AI training. In a joint study with *MIT Sloan Management Review*, we found broad consensus among global leaders across 29 industries that humans need to build and maintain expertise to verify AI solutions and outputs. These leaders agree that the most significant AI risk is cognitive offloading that leads to outcomes that are misaligned with public values or that harm vulnerable groups.

Establishing Accountability and Continuous Learning at the Top. Someone needs to be accountable for ensuring the AI governance system is effective and scalable to keep up with continuously evolving risk and opportunities. The goal is to have a senior executive build and oversee a mature three-lines-of-defense model that divides responsibilities among frontline operations, risk management, and monitoring and auditing. The model should also be implemented across the organization to prevent a diffusion of accountability and passing the buck.

Executive accountability for enterprise-level decisions about AI is a table-stakes best practice. A majority (80%) of the 142 respondents reported that accountability for enterprise-level AI decisions is managed by named executives or a standing committee, though only a minority (less than 10%) claimed it's a single-named accountable executive. Yet, even when accountability for enterprise-level AI decisions is shared, when asked who holds primary accountability for AI outcomes and risks in the organization, 94% were able to select a single executive.

It remains to be determined who the accountable executive should be: 75% of organizations choose among the chief AI officer, chief information officer, and chief technology officer; only a minority chose the chief risk officer. Regardless, this individual should be supported by a cross-functional committee with the expertise to evaluate decisions from every functional angle because of the breadth of AI impact. The committee needs to be able to determine how decisions should be made, scope out what new roles are required for effective oversight, and assess how to allow for situational nuances that balance risk and innovation.

Infusing the Organization Chart with the Necessary Accountable Roles. Accountability should cascade across the three lines of defense to include business and functional leaders and end users. A shared governance approach is especially critical since outcomes from AI workflows can be produced at a scale and speed no human review committee can match.

Cascading governance to the end user is a challenge for most organizations. Only 52% of survey respondents said their accountability for AI outcomes extends to end users. The majority of those respondents are with accountable innovators or cautious adopters.

Deciding who should be the accountable owner in any scenario is context specific. For example, if a tool has been designed with the expectation that the end user accepts the output (for instance, an answer to a first-line IT question), then the central function that built the tool may be responsible for its outputs. However, if a tool is built for general-purpose use, it may require the end user to have final accountability.

At a high level, every decision should be made by the individuals with the highest level of visibility, capabilities, and control over the situation. Every decision owner then needs to be clearly documented, so they are easy to identify when a question or problem arises. AI governance also introduces roles for activities that did not previously exist at this scale: red teaming new builds, assessing drift, revalidating when an AI model version changes, as well as monitoring and auditing agent behavior in production.

To prevent the diffusion of responsibility, every line of defense needs a clearly named owner with clearly documented handoffs. At BCG, we assign clear ownership spanning AI product builds, maintenance, and risk compliance. We also create new roles to monitor internally built AI tools for compliance and drift. We then ensure everyone can build expertise in their role. For example, the person accountable for ensuring compliance with the General Data Protection Regulation in every AI implementation needs to have sufficient expertise and experience to review any AI tool effectively.

Creating a Unified, Risk-Tiered Governance Framework. Every organization should define a formal risk-tiering framework, supported by triage and classification, to identify which use cases are a low risk (no review required), moderate risk (some review required), high risk (significant review required), and critical risk (no-fly zone). This formal classification process ensures that employees and decision makers have a consistent basis for determining what can proceed, what requires additional scrutiny, and what should not be pursued. Accountable innovators do not develop these frameworks in isolation. In addition to their own enterprise standards, they reference, align to, or are certified by multiple external governance frameworks, such as the EU Artificial Intelligence Act, the ISO and International Electrotechnical Commission 42001 standard, and the National Institute of Standards and Technology AI Risk Management Framework.

Still, because AI governance touches various functions, including IT, legal, human resources, data protection, and cybersecurity—each of which addresses different risks and often with different terminology—organizations may run into conflicts across policies; more than one-third of our respondents reported having a conflict. However, 39% of respondents said they have specified a committee to make the call—likely because there was no legal requirement involved in the situation.

When designing a risk-tiering framework, leaders should include a way to systematically document every decision alongside its underlying analysis and assumptions. This documentation can then

be used to automate use case triage and risk classification to scale decision making. Surprisingly, only 50% of respondents claimed to do this. In addition, only 26% said that when a novel AI question is resolved, the answer is systemically incorporated into standing guidance for others to follow.

The idea is to turn a policy library into a working system that allows for continuously improved decision making. Organizations that both record and codify AI decisions are twice as likely to have more than 50 use cases in production compared with those that don't take these steps. Organizations that record and codify also have more than 30% higher ROI compared with those that only record. They will also likely be prepared to engage with regulators; emerging regulations are largely built around risk tiers, documented human oversight, and evidence of reasonable steps to reduce risk.

Streamlining and Automating Processes and Controls Where Possible. When every employee in an organization can build an agentic tool, leaders need to make it easy for them to comply with policies. However, doing so can put a system of governance under strain when there is an unprecedented volume of requests and emergent risks, since each request often requires numerous decisions made by multiple people. (See “The Challenges of Implementing AI Governance.”) Most AI governance systems can promptly manage a handful of requests, but response time slows when thousands of employees are innovating and sharing ideas.

– The Challenges of Implementing AI Governance

Almost every organization is discussing how to balance risk with AI innovation. In our survey, we included a free-response question, asking for leaders’ single biggest struggle in making AI governance operational. We found several recurring challenges.

Challenge	Respondents (%)	Description	Example Quote
Enablement (training, culture, and buy-in)	19.0	Training, awareness, skills, and building a culture and buy-in for AI	“Getting an organization of our size to attend the trainings.”

Difficulty making clear policies	14.1	Hard to write clear, consistent, usable governance policies	“Creating a consistent framework that is widely communicated and understood.”
Speed of AI development	12.7	External AI tech evolves too fast to keep up with	“The velocity of new AI tool introduction makes it almost impossible to keep up.”
Ownership and accountability	9.9	Unclear who owns decisions and is accountable	“No clear ownership—no one accountable for the big picture.”
Speed of enterprise adoption	9.9	Own organization is slow to roll out and approve new tools	“Rollout to enterprise level was slow; bureaucracy of process.”
Balancing innovation versus control	7.7	Tension between innovation speed and governance and controls	“Tension between speed imperatives and required controls.”
Regulatory burden	6.3	External regulation and compliance burden	“Complex regulatory and compliance frameworks that control

			bank operations.”
Resource constraint	5.6	Not enough time, budget, or capacity	“Budget, resources, volume of requests, and the speed at which it is growing.”
Fragmentation and coordination across functions	4.9	Silos and poor coordination across functions and regions; duplicated work	“Siloed out functions and regional silos.”
Shadow AI	4.9	Untracked AI use with no visibility into the data	“Shadow AI is prevalent and impossible to track.”

Experts in our interviews elaborated on their challenges in more detail, citing the following.

Policy Without Process. A thorough, well-drafted policy that nobody has translated into a process with which people can actually make decisions is indistinguishable, operationally, from no policy at all.

Generic Policy Design. Guidance that is written to cover everything tends to miss where the day-to-day risk sits. Policies come out either too abstract to apply to a specific use case or so detailed that they lose the big picture. A director in banking described the problem as such: assigning roles and responsibilities across a shared matrix makes AI accountability tricky, because every area has its own use cases and rules that cannot be applied generically across the board.

Manual Processes. Even organizations that have streamlined their governance are usually still running it manually, because building the automation is real engineering work that competes with business priorities. This is survivable with a low volume and impossible at scale.

Governance Overload. Without a clear risk taxonomy and a named decider, the same use case gets reviewed by cybersecurity, legal, and privacy functions in parallel, each reaching a different opinion, with nobody empowered to close it out. One leader described the result as endless churn. About 20% of our survey respondents also named speed of decision making and unclear ownership as their single biggest struggle.

Missing Roles. The processes assume continuous work—monitoring drift, revalidating after model updates, monitoring and auditing agent behavior, and reteaming new builds—that is not in anyone’s job description. Organizations design the workflow and then discover there is no one to run it.

Fear of Creating Performative Compliance. Where mistakes are punished, people stop surfacing them. Worse, they socially perform innovation instead of being genuinely innovative: the mandated pilot gets run, the usage gets logged, and nothing about the work changes. A vice president in insurance said that the challenge extends to getting employees to be transparent, because they are afraid that the models and tools they are already using will not be approved.

Genuine Limits on Transparency. A lack of transparency is not always a design failure. True monitoring and the ability to audit multiagent systems and individual usage is technically hard and has to be implemented in compliance with privacy laws. Leaders can monitor the tasks their team runs through a compliance application programming interface, but they cannot easily see when someone builds a skill or dashboard and shares it with coworkers. Leaders described this as a whack-a-mole scenario, and one recalled finding people installing tools and plugging personal machines into the network. Shadow AI, in the words of one health care vice president, is prevalent and impossible to track.

Accountable innovators create a single place to submit requests to make compliance easier, streamline or automate reviews, and speed approvals of low-risk innovations. Our survey data suggests that companies implementing this approach are 30% more likely to approve the average use case within four weeks and reduce shadow AI by 15%.


For example, one organization found that accountable leaders were overwhelmed with approval requests for innovative AI use cases. The result was that it took more than two months to make simple decisions. To solve this, leaders created a risk hierarchy and automated the obvious decisions that had historical precedent. This streamlined decision making and approvals for new AI citizen-developed tools—and 80% of requests were approved within hours.

Governing all the emerging risks at this scale also requires technical controls. Some examples of technical controls include: ensuring every agentic tool (such as Claude Code or Codex) runs in a controlled way, installing compliance skills to guide the AI’s behavior on every account, and running automated monitoring using a compliance applications programming interface similar to that used for data-protection monitoring. Both human and platform controls should also become more stringent as risk increases, and organizations need to ensure they are actively used in practice. High-accountability organizations are more than two-and-a-half times as likely as others to have and deploy a kill switch—the ability to quickly shut down an AI system that is already in production.

BCG’s 2026 Applied AI survey shows that an emerging best practice among companies with mature governance is to build out an organization-wide AI platform to increase transparency and control of all agentic workflows without limiting what teams can build. These organizations are also two times as likely to be embedding agentic AI into their legacy modernization strategy, compared with those that have less mature governance, and more than three times as likely to be implementing agents that strategically drive their modernization. (See Exhibit 5.)

EXHIBIT 5

High-Governance Organizations Are Pursuing Central AI Platforms and Agentic Modernization



Source: BCG’s 2026 Applied AI Index.
Note: n = 1,330. High governance is a composite measure of responsible AI (enterprise-wide RAI policy, enforced controls, and board-level accountability); data governance (clear data ownership, automated controls, and stewardship enforcement); data compliance and protection (continuous privacy monitoring and automated compliance); cybersecurity (zero-trust architecture, continuous monitoring, and rapid incidence response); and agentic governance (implementation of six agentic governance controls).

Most leaders are engaging in active discussions about AI governance but finding it incredibly challenging to implement in practice. AI governance is never finished—the technology changes constantly, new use cases arrive with each evolution, and regulation is evolving. Yet, our research suggests that AI governance and innovation go hand in hand to achieve tangible ROI.

Addressing the AI governance challenge requires accountability at every level and systemic transparency. Leaders need to know what is happening in their organization, make sure someone is accountable for every substantive decision and process, and build in a scheduled review of their own model. Leaders should also know what needs to be escalated, what is approved

automatically that should not be, and where the process is bypassed. Moreover, leaders should treat these signals as the input for governance and process revisions—rather than as failures.

The authors thank Maarten Beckx, Miriam Feldman, Shekhar Iyer, and Harsh Kumar for their substantial effort conducting this research, as well as Clemens Nopp for his contributions of data, analysis, and expertise.



BCG Institute is Boston Consulting Group's real-world think tank. We research the future of thinking about the future. Built on BCG's decades-long intellectual legacy and a global network of academics, scientists, and institutional partners, we bring together expertise across disciplines to advance ideas before they become conventional wisdom. Tracking the forces reshaping business, technology, economics, and geopolitics, we turn emerging patterns into clarity, helping leaders around the world turn early insight into measurable outcomes.

Authors



Matthew Kropp

Managing Director & Senior Partner; BCG Institute Fellow
San Francisco - Bay Area



Julie Bedard

Managing Director & Partner; BCG Institute Fellow
Boston



Ulrike Schwarz-Runer

Managing Director & Senior Partner; General Counsel
London



Steven Mills

Managing Director & Partner, Chief AI Ethics Officer, Global Leader, BCG Center for Digital Government
Washington, DC



Anne Kleppe

Managing Director & Partner
Berlin



Lisa Krayer

Partner
Philadelphia



Marc Roman Franke

Partner & Associate Director, AI and digital transformation
Berlin



Zeeshan Shah

Director, BCG Platinion
Nashville



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.

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.

© Boston Consulting Group. All rights reserved.