



ARTIFICIAL INTELLIGENCE

The Agentic Leadership Playbook: A Scaling Strategy for CTOs and CIOs

A conversation with BCG’s Mark Abraham and Neveen Awad

ARTICLE JULY 08, 2026 12 MIN READ

Two practitioners who’ve helped shape and implement agentic AI strategy for some of the world’s largest companies share what’s working—and what’s stalling progress.

Everyone’s talking about agentic AI. What are CTOs, CIOs, and their teams asking when they first walk in the room?

Mark Abraham: First, there is a lot of confusion about agentic AI: where to deploy agents, how autonomous they should be, and what’s real versus hype. Second, business unit leaders want to address cost pressures. They are looking at agents as a way to hit productivity targets, whether that’s getting cost-per-asset down or finding headcount efficiencies. But what is often missed is that this is a three-part game. Speed first, growth second, and cost third—in that order. By setting the right agentic AI strategy, the boldest companies are already tripling campaign speed, tripling ROI on media budgets, and taking 15% to 20% out of total functional spend.

Neveen Awad: The starting question is the same for any transformation: Where is the value? Agents are an incredible technology and an expensive one. If a process is simple and rule-based, you don’t need an agent. Agents work best in situations with complexity—multiparty information exchange where interpretation is needed. That’s where an agent earns its keep. Procurement is a good example. Figuring out when and how much to renegotiate with suppliers, initiating that dialogue. Agents can do most of that, with humans handling only the final conversation.

When it comes to agentic AI implementation, what assumptions do clients most often get wrong?

Neveen: At one end, you have the “agents are going to do everything, let’s go” camp. At the other end are clients who have stopped thinking about this as a technology play entirely. For them, it’s a business orchestration play. But the point is to use technology as the mechanism for redesigning how the whole organization thinks and works. The leaders don’t ask, “How do I deploy agents?” They ask, “How do we work differently?” And they’ve accepted that perfectionism is a liability now. You pilot, you measure, you iterate.

Mark: Sixty percent of tech stack and agentic deployment decisions now sit with the CIO or CTO. That's a real shift from even the last AI cycle, where business had as much or more of the say. The danger is that it becomes tech-led rather than a true business-tech partnership. The farsighted CIOs insist on clarity from business counterparts first. They ask, "What are your goals, your strategies, your priorities?" Then they work their way back to the technology requirements.

How should CTOs and CIOs think about data readiness as part of their agentic AI strategy?

Mark: The stakes have gone up on data readiness because agents aren't just generating recommendations, they're taking action. A bad definition of "lapsed customer" can cause an entire campaign to go awry. At the same time, agents can now collapse what used to be an eight-week data gathering and cleaning process into a single week. What remains irreplaceable, though—and it's the most underappreciated part of deploying agentic solutions—is what I call the intelligence layer. By that I mean mapping your actual business context onto the data, so agents understand the drivers of your business.

Neveen: What agents need isn't perfect data, it's honest data—data you trust. CTOs and CIOs should build and test agents against that core, see how they're reasoning and making inferences, and then expand the data they're working from as confidence grows. And don't assume GenAI will just clean the data for you, so you don't need to worry about it. There is still significant work to do on the underlying architecture. The agents can be far more accurate and efficient when you've pruned the context they're working from.

Say more about that intelligence layer. What role does it play in agentic AI architecture?

Mark: Every company has its own theory of the business—the KPIs it tracks, how those KPIs link together, the micro-segmentation that matters. For one of our retail clients, that means thinking about existing versus new customer traffic, seven specific segments, micro-geographic demand pockets, price competitiveness signals. When you map that onto the data, any agent you put on top, even an off-the-shelf insight tool, suddenly has the context within which your team thinks about performance. That’s how you get around hallucinations and inconsistent answers. We believe large companies should own this intelligence layer. Our BCG X team holds a patent on an approach to building this efficiently, which we call EnterpriseIQ. You can plug almost any agent workflow into it and get reliable results. And because the intelligence layer is separate from the agents themselves, you can swap models in and out as better ones emerge.

Neveen: I think of it as smart plumbing. You have data over here, people making decisions over there, systems elsewhere, and organizations are constantly synthesizing all of that to take action. The goal is to build the plumbing that gives agents enough context to make those same decisions, and eventually better ones. You’ve always had APIs. What’s new is the importance of context, and all the work that goes into context engineering: what’s relevant and what isn’t. The shortest path is to start with a specific problem you’re trying to solve. Build the knowledge graph around that. Get the agents working against it. Then look at where the answers are breaking down—is it the agent logic, the knowledge graph, the source data—and go up and down the stream until it’s right. Then expand.

Enterprises run on their ERPs and CRMs. How does an agentic AI strategy change what CTOs and CIOs should ask of those systems?

Neveen: You want to separate agentic logic from core platforms entirely. Keep the core system minimalist and as close to out-of-the-box as possible and move all the custom logic into the agentic layer. That reduces your tech debt significantly, because your core platforms become easily upgradable. You may also find you don’t need the top-tier, bells-and-whistles version of a system anymore because agents are handling the complexity that used to require it. Above the core platforms, you need a model gateway to function as a unified enterprise AI orchestration layer and the knowledge layer that handles shared context and memory. The key is modularity. You should be able to swap LLMs in and out as leading providers ship better versions, without having to rebuild your agent logic.

Mark: I'd add the layer that I don't see enough companies invest in is a marketer-facing UI, a single pane of glass that sits above all of it. The technology is moving so fast that you don't want to train people on specific tools that will change in six months anyway. A well-designed interface abstracts all of that and makes the stack composable without burdening users with its complexity. It also becomes your single point of control and security. You can program different access levels, different authorities for different user types, and make agentic AI governance practical.

Enterprise AI governance is not new. What does agentic AI governance require once agents start acting on their own?

Neveen: The answer is graduated autonomy: shadow mode, supervised mode, guided autonomy, full autonomy. Each tier is earned through demonstrated performance, each requiring clear measurement of the outcome you're shooting for and the risks you're trying to mitigate. The harder ongoing work is that every time you upgrade a model, agent behavior will likely change. You need continuous evaluation pipelines, and you need to require agents to report their journey. Many leaders are also standing up red teams whose job is to make agents behave badly before you find out at scale.

Mark: You have to define the guardrails with the business before you start. If a marketing function doesn't have clearly defined brand standards, you cannot have a conversation about autonomous campaigns. When we built a briefing agent for a beauty company, we found it was agreeing with users and flattering their initial drafts rather than pushing back on weak briefs. We reprogrammed it to be "Socratic," to always ask the ten questions that the best marketers would subconsciously ask themselves when writing a brief. Companies are competing not just on data but on their ability to codify the best practices of their best talent into the tools.

The BCG AI Radar data shows a stark gap between executives

who think agents will deliver real value this year and those who believe true transformation is still three-plus years away. What's your read?

Neveen: What struck me most in that report wasn't the transformation timeline. It was the gap between how much the executive level was being trained and engaging with GenAI and how much the working level was. Transformation happens where the work happens. The only way you can truly enable it is to go value stream by value stream and really think about what it would look like to reimagine each one. That will create massive business transformation, but the people who own those end-to-end processes have to be a big part of the change. The trailblazer organizations are also the ones who had trained more than half their workforce. Agentic AI adoption and transformation are the same track.

Mark: Only 5% of companies are truly agent-first. And it's not a spectrum, it's two clusters. The AI-first companies have driven adoption above 80% of basic LLM tools across their organizations. The other 95% haven't given employees the latitude to even experiment. Agentic AI adoption rates sit at around 30% or less for even the basic tools and are often patchy at that. What separates the two clusters comes down to co-creation. The companies at the forefront build their agentic solutions with their best people around their real day-to-day challenges. A media company I work with took five members from different entertainment teams and involved them in shaping what the agent solution should look like. They iterated on these versions weekly until they could react quickly with targeted campaigns. From there, the company changed its processes and built an adoption approach that scaled to thousands of teams.

What kills momentum in moving agentic AI implementation from pilot to scale?

Neveen: The momentum killer is forgetting that 70% of getting to scale is people and change, 20% is data and technology, and only 10% is the algorithms. Top-down leadership has to really believe

and say we're doing this, it's going to be hard, there will be waste, but we're doing it. People want to be part of winning teams. Start small, show real change, celebrate it, and others want in. The second killer is giving up too early when getting agents to communicate and share context gets hard. When a multi-agent network hits a context gap, that's not a failure. It's just something to build. The third is agentic AI governance chaos—unclear ownership creates a lot of “I own that, no I own that” confusion that stops everything.

Mark: I'd add: not moving from pilot to scale fast enough once the initial deployment works. There's a massive difference between a use case that works in an isolated environment and one that reshapes a function. Scaling involves changing incentives, changing processes, serious training—and it requires sustainment. The C-suite churns a lot. New leaders come in with new visions and don't build on what was learned. If the strategy only lives in the heads of the current leadership team, it's not a strategy, it's a bet on continuity. Boards are increasingly demanding a real AI roadmap with a portfolio view, and that accountability at the board level is one of the things that protects against momentum loss.

What do you most want CTOs and CIOs to take away from this?

Neveen: Value-driven intellectual honesty. We are in a period where there is an enormous amount we don't know. The leaders who acknowledge that openly—who say, “Here's what we think, here's what we're going to try, here's how we'll know if it's working”—will build the trust that makes these transformations succeed.

Mark: When it comes to agentic AI, bring together your best people across the relevant functions from day one. Let them shape the solutions and define what “good” looks like. Ways of working are critical and more than just agile. The use cases have to be big enough to unlock real value but small enough that nimble teams across business, tech, and data can deliver.

Featured Experts



Mark Abraham

Managing Director & Senior Partner; Global Leader, Marketing, Sales & Pricing Practice
Seattle



Neveen Awad

Managing Director & Senior Partner
Detroit



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.