VirtuousAI Chief Executive

State of AI Adoption in the Mid-Market

How CEOs are approaching and using AI

For most CEOs, the conversation around AI has shifted from curiosity to pressure. Board members are asking about it. Employees are experimenting with it. Competitors are talking about it—sometimes loudly. And yet, behind the scenes, many leadership teams are still trying to answer a much more basic question:

What, exactly, should we be doing with AI right now?

VirtuousAI | Chief Executive

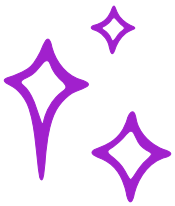

That uncertainty is especially pronounced in the mid-market. Unlike large enterprises, mid-sized companies don't have the luxury of massive data science teams or unlimited budgets to test and discard new technologies. At the same time, they can't afford to sit still. AI is moving quickly and the cost of falling behind feels real—even if the path forward does not.

To understand how CEOs are navigating this moment, **Chief Executive** partnered with **VirtuousAI** to survey mid-market leaders on how they're using AI today, where they see the greatest opportunity, and what is preventing broader adoption. The findings paint a clear picture: belief in AI's potential is nearly universal, with nearly all CEOs (98.5%) seeing value in AI for their business.

Execution, on the other hand, remains tentative. Most companies are experimenting at the edges and focused on efficiency gains rather than transformation, and still struggling to move from isolated use cases to enterprise-wide impact.

The following report is a look at where mid-market AI adoption truly stands today, the gaps between ambition and action, and what CEOs actually need in order to close them.

Key Findings

In Fall 2025, on behalf of VirtuousAI, Chief Executive asked over 300 CEOs around the country to reflect on where they were on their company's AI adoption journey and what challenges were holding them back.

- **Nearly all mid-market CEOs (98.5%)** say AI has value for their business, with only 1.5% saying they do not see a role for AI.
- **Just 7 percent** of CEOs report having a company-wide AI strategy with multiple initiatives in place, while 52% say they are still in the pilot phase and 31% have explored AI but not implemented it.
- **Seventy-eight percent** of CEOs cite efficiency and cost reduction as the primary objective of their AI initiatives, compared with 19% who say developing new products, services or revenue streams is a top priority.
- **More than half (53%)** say they want to use AI to improve data analysis and decision-making, signaling that most AI efforts remain focused on insight rather than execution.
- **Eighty-six percent** of CEOs identify lack of AI expertise as a major barrier, while 81% cite difficulty integrating AI with existing systems; 65% point to data quality and accessibility issues.
- **Six in 10 CEOs** report having active AI projects underway, despite the lack of enterprise-wide strategies.

98%

of CEOs see value in AI for their business

AI's Value Is No Longer in Question

Nearly every CEO surveyed agrees that AI has a role to play in their organization. In a business landscape where most emerging technologies face skepticism in at least some quarters, this kind of consensus is rare. When asked about which AI applications held the greatest potential, respondents pointed most often to three areas: generative AI (60%), process automation (60%), and predictive AI (51%).

Belief in AI's potential, however, has arrived faster than clarity. While the majority of CEOs appear convinced that AI will transform their businesses, many are still unsure where it belongs operationally, what problems it should solve first or how to move from general enthusiasm to specific action.

As a result, AI often enters the organization through experimentation—teams testing tools, leaders greenlighting pilots, marketing departments deploying chatbots—rather than through a top-down business strategy.

Most Companies Are Experimenting—Not Executing

The survey data shows a stark gap between intent and action. Only 7% of CEOs say their company has a coordinated AI strategy with multiple initiatives in place. By contrast, 52% remain in pilot mode and 31% have explored AI without implementing it meaningfully.

This matters because isolated pilots don't compound. A team tests ChatGPT for customer service, sees some productivity gains, and calls it a win. Marketing launches a different AI tool. Finance explores another. Six months later, that company has multiple disconnected experiments running in silos, generating insights that never talk to each other, creating what technologists call technical debt. The company has spent money and time but hasn't actually learned anything that carries forward.

7%

Just 7% of CEOs say they have a company-wide AI strategy, while over 80% report experimenting in the early stages with pilots or tactical applications.

The problem is structural. Without shared data, clear ownership, and defined success metrics tied to business outcomes, pilots tend to stall. They stay interesting side projects rather than becoming strategic capabilities. The intelligence created in one department doesn't feed into decisions elsewhere; there's no roadmap for moving from experimentation to production and no connection between initial insights and broader operational impact.

The result is AI investment that produces incremental gains in isolated functions but

does not transform how the business actually runs.

Efficiency Dominates AI Priorities—for Now

When CEOs were asked where they expect AI to deliver value, nearly four in five respondents (78%) cited cost reduction and productivity gains as their primary objective, while just 19% pointed to new products, services or revenue streams.

This focus reflects both caution and pragmatism. Efficiency gains are easier to measure, easier to justify to the board, and easier to pilot without disrupting core operations. Automating a routine process, for example, requires far less organizational change than building a new business model around AI.

78%

of CEOs are seeking to improve efficiency or reduce costs.

Only 19% are looking at AI as a transformational opportunity.

The efficiency-first approach also reflects something else: most CEOs are still treating AI as a tool to improve what already exists, not as a foundation to build something new. A company might bolt an AI extension onto its CRM system and that may improve productivity, but it won't be game-changing. These individual solutions also risk creating more silos and, ultimately, more technical debt. CEOs who want AI to drive growth need to connect early efficiency wins to broader strategic outcomes.

But there's also a cost to this framing. If you use AI to make your current way of doing business 10% faster or cheaper, you've made progress. But if your competitors use AI to reimagine what the business does entirely—to create new revenue streams, serve new customers or unlock margin in ways the old business model didn't allow—

then you've fallen behind.

Imagine... a logistics provider that uses Generative AI to automate RFP responses, cutting drafting time by 50%. That is the efficiency win. But because they integrated that AI agent directly with their dynamic pricing model, the system doesn't just write proposals faster; it intelligently adjusts rates in real-time based on predicted route capacity and margin targets. The result isn't just time saved—it is a strategic filter that ensures the sales team wins profitable business rather than just more business.

More than half of CEOs (54%) say they want to use AI to improve analysis and decision-making. That still positions AI primarily as an intelligence layer rather than an execution layer.

So while leaders are comfortable asking AI to surface insights, to analyze data and tell them what it sees, they are less comfortable letting AI trigger actions. The trust issue is understandable, to be sure. If AI is telling you something counterintuitive, you want to understand why before acting on it—and if it recommends a course of action, you want to verify the logic before committing resources.

But the longer AI remains upstream of operations, the harder it becomes to realize measurable ROI. A system that analyzes your customer data and tells you that high-value customers are churning at an unexpected rate has value. But a system that not only identifies that pattern but automatically takes action—e.g., triggers a retention offer, adjusts pricing, prioritizes support—delivers an order of magnitude more value.

The data suggests that most mid-market companies haven't yet crossed that threshold from AI-as-analyst to AI-as-operator. They're building an intelligence layer, but it's not *connected intelligence*. The acceleration comes only when intelligence is connected to execution.

What's Holding Back Progress

Eighty-six percent of CEOs identified lack of expertise and skilled talent as a top challenge, with 81% citing difficulty integrating AI with existing systems. Nearly two-thirds (65%) point to data quality and accessibility issues.

86%

of CEOs say a lack of internal expertise and skilled talent is a top challenge in their AI adoption journey.

These challenges are mutually reinforcing. Fragmented data undermines trust in AI outputs. Lack of expertise slows integration efforts. Poor integration increases technical debt and makes future projects harder. Together, they widen what many CEOs experience as an execution gap—or the sense that AI's capabilities are advancing faster than their organization's ability to absorb and deploy them.

These are solvable problems. But they require a different approach than most organizations are taking. Buying more AI tools, for example, won't advance execution because an organization with fragmented data, dispersed expertise, and siloed systems doesn't need another point solution or tech tool. It needs integration and a framework that connects data, decision-making and action.

Bridging the Execution Gap: From Tactical AI to Strategic AI

The survey data points to a fault line between two very different approaches: tactical AI and strategic AI.

Tactical AI shows up as:

- **General-purpose tools** (ChatGPT, Gemini, Copilot) that can boost individual productivity but aren't trained on your business
- **Bolt-on copilots** that extend existing systems like CRM or Microsoft products, giving you insights from that silo but not the full picture
- **Single-function agents** deployed for a specific purpose—customer service, HR—that solve one problem but often deepen data silos

Each of these can deliver value. A chatbot can save your team time. A CRM copilot can surface customer patterns. A recruitment agent can screen résumés faster. But each operates on its own, fragmented from the rest of the business and, therefore, each adds another layer to an already siloed technology stack.

Strategic AI, by contrast:

- **Ingests everything** — your data across functions, internal and external, structured and unstructured, unified into a single view
- **Classifies signals** — distinguishes meaningful information from noise so decision-makers can see what actually matters
- **Drives business outcomes** — is wired to automate actions at the operational level and move specific KPIs

Most mid-market companies are still operating tactically—not by choice, but because it's the path of least resistance. A CEO can approve a pilot with a general-purpose tool in an afternoon. Standing up a unified data layer that connects sales, operations, finance and marketing is harder, requiring architectural work, governance and change management.

Yet that shift to strategic AI is where compounding value lies. That is where AI stops being a collection of clever tools and starts becoming part of how the business runs.

THE VIRTUOUS CYCLE | WALK > RUN > SCALE

Understanding the **Virtuous Cycle**

Strategic AI is not just a different kind of software; it is a different operating model. At its core is a simple but powerful idea: a virtuous cycle of learning and action. Every action makes the next action smarter. Every decision makes the system stronger. Every cycle drives real business outcomes.

That is what separates a strategic approach from a purely tactical one. When AI is used only to make a process 10% faster or cheaper, the organization wins a local victory. When AI is used to build a system where every outcome feeds back into decision-making—where, for example, each transaction teaches the system how to optimize for the company's core KPIs—the effect compounds.

Consider a supply chain example. Tactical AI might apply machine learning to forecast demand more accurately. Helpful, but narrow. Strategic AI would connect that forecast to purchasing, pricing, marketing and inventory, adjusting all of them in concert to optimize for margin and cash flow. Every time the system reorders or a promotion runs, it learns. The model becomes sharper with each pass.

To make that kind of virtuous cycle real inside a mid-market company, you need a way to build it step by step. That's where the Walk-Run-Scale framework comes in. It breaks strategic AI into a sequence of capabilities—first seeing clearly, then acting faster, and finally, optimizing continuously across the business.

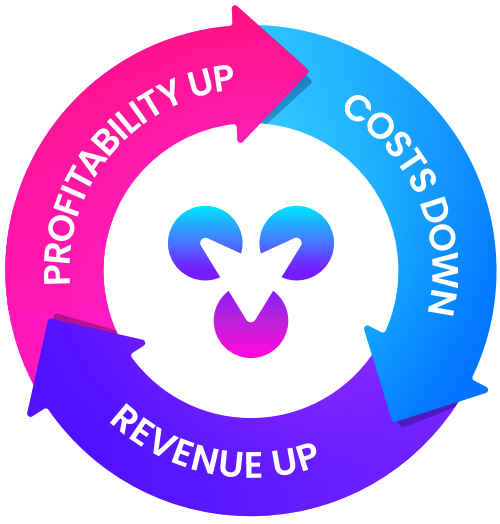

THE VIRTUOUS CYCLE **WALK** > RUN > SCALE

Phase 1: **Walk**
The Connected Intelligence Phase

Most mid-market companies have the data they need, but it's scattered. A customer is known in the CRM, but the support history lives in a different system. Inventory is tracked in the ERP, but sales happen in Shopify. Returns data sits in another system.

The Walk phase brings all that together. You integrate your key data sources—ERP, CRM, customer support, inventory management—and create a unified view. For the first time, the system can see correlations across the business. It can understand relationships between functions that were previously invisible because the data wasn't talking.

The outcome: strategic intelligence. Not just 'We sold 100 units' but 'We sold 100 units, returned 15 due to sizing, and those 15 were disproportionately in sizes 34-36 of our premium line.' That's intelligence you can act on.

THE VIRTUOUS CYCLE WALK > **RUN** > SCALE

Phase 2: **Run**
The Automation Phase

In this phase, you deploy automation agents that take discrete actions based on the patterns the system has learned. Initially, this happens with what's called 'human-in-the-loop': the system makes a recommendation, a human approves it, and the action happens. This builds trust. Leaders get comfortable with the system's judgment. They see that the recommendations work.

Once trust is established, you can automate further. The human step becomes optional. The system has learned enough that it can execute high-confidence scenarios on its own.

The outcome: operational efficiency and speed. Your high-value people stop doing repetitive decision-making, and your response times drop from days to minutes.

THE VIRTUOUS CYCLE WALK > RUN > **SCALE**

Phase 3: **Scale**
The Optimization Phase

Now you're no longer managing isolated automations, but orchestrating a set of automations that work together, continuously optimizing for your strategic outcomes. Maybe that's optimizing inventory while simultaneously managing pricing and marketing. Maybe it's balancing supply chain efficiency with cash flow. Maybe it's coordinating customer acquisition, retention and lifetime value.

Most critically, the system sees all of it, and optimizes across all of it. And it learns from every transaction, every decision, every cycle.

The outcome: strategic advantage and margin expansion. Your business is running smarter, you're capturing opportunities you couldn't see before. And you're avoiding problems you couldn't predict previously.

Mugsy

CASE STUDY:

Starting With the Outcome

Mugsy is a fast-growing men's apparel brand known for premium jeans and comfort-focused fits. As the company scaled, a familiar problem emerged: inventory had become too complex to manage by gut and spreadsheets alone.

Customer, sales, inventory and returns data lived in different places. Shopify handled e-commerce. A separate returns platform captured send-backs. Customer service ran through its own system. Third-party logistics providers and internal tools tracked warehouse activity. Over time, Mugsy added "20 or 30 plugins" to the website, each with its own numbers.

> If I asked a relatively simple question like,
> **'What was our return on ad spend yesterday?'**
> I could get four different answers
>
> – Leo Tropeano, Mugsy CEO –

CASE STUDY Mugsy

At the same time, Mugsy had become an extremely SKU-heavy business. What began as a handful of jean styles expanded into thousands of SKUs across sizes, colors, and seasonal lines. Hitting what the team considered a strong performance—selling through roughly 70% of inventory—still meant being left with a long tail of slow-moving product tying up cash and eventually pushed out via discounting.

Rather than start by shopping for another forecasting tool, Tropeano started by defining the outcome he wanted: better inventory and data decisions that would free up capital, reduce unnecessary discounting, and give him a clearer, trusted view of the business day to day. From there, the work followed a structured process.

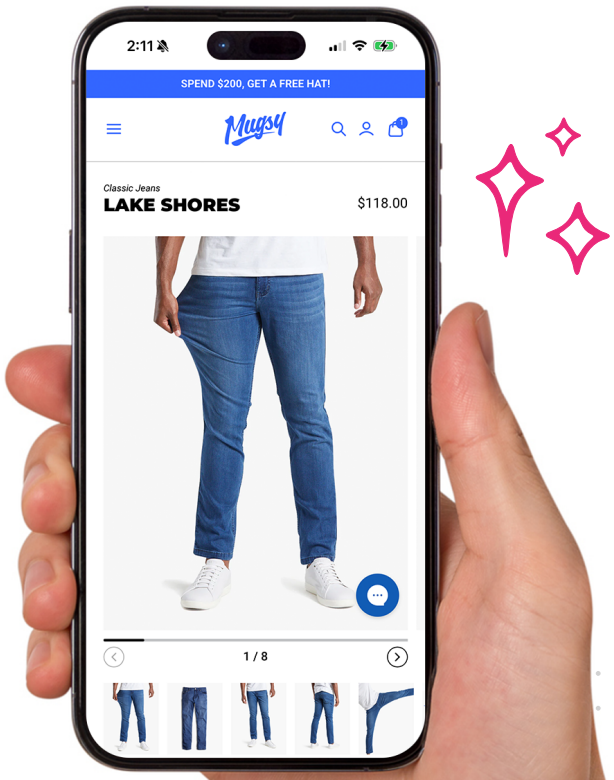

To bridge the gap between their ambitious goals and their fragmented data, Mugsy partnered with VirtuousAI. By adopting BAIO (Business Automation, Intelligence & Outcomes), Mugsy moved beyond the limitations of isolated point solutions and committed to a strategic framework capable of unifying their entire data ecosystem.

CASE STUDY Mugsy

STEP 1: Clearly define the problem in business terms

For Mugsy, that meant recognizing that they had a large amount of working capital trapped in inventory and an increasingly unreliable picture of what was driving performance. Inventory planning based on manual judgment and inconsistent data could no longer keep up with the growth and SKU complexity of the business.

STEP 2: Define the assets

Next came mapping the people, systems and data involved. On the people side, that included Mugsy's CEO, COO, merchandising and inventory-planning leaders. On the systems side, it meant Shopify, the returns platform, customer-support tools, warehouse and third-party logistics systems, internal spreadsheets, and other plugins that contained relevant signals about demand and behavior. The picture was fragmented, but the raw material for better decisions already existed inside the company.

STEP 3: Build the foundation

With the assets identified, the team turned to building a foundation: a model of how the business works that an AI system could actually reason over. In practice, that meant defining the core concepts—customers, orders, products, vendors, channels, inventory locations, price points, seasons—and the relationships among them.

This "corporate brain" work did not involve writing a forecasting algorithm from scratch. It involved encoding Mugsy's own semantics and KPIs into a structure that could later be trained on real data—e.g., what constitutes an order, how returns relate to margin, how marketing activity affects sell-through by channel, and so on.

CASE STUDY Mugsy

STEP 4: **Load and train the platform**

Only after this framework was in place did the team begin loading historical data. Sales, returns, inventory positions, and marketing performance were ingested so the system could start to learn patterns, such as which products move quickly, which have higher return rates, and how seasonality and promotion influence demand.

Earlier experiments with other AI tools had made Tropeano wary; he had seen systems that could not reliably answer basic math questions, which reinforced his view that serious AI work requires humans who understand both the algorithms and the business context. In this phase, the focus has been on getting to the point where leadership could ask, "What were sales yesterday?" or "How did each store perform?" and receive fast, consistent, trustworthy answers from a single source.

STEP 5: **Apply the use case (Walk > Run > Scale)**

Today, Mugsy is in the Walk phase, which means the company is using the unified data and "corporate brain" to answer questions that used to require stitching together multiple reports: daily sales, performance by product and channel, emerging patterns in returns and margin. Tropeano describes this as having a "ChatGPT for Mugsy" at his fingertips—an interface that lets him live in the numbers without spending hours pulling them.

In parallel, the team has developed an AI-driven inventory planning model that tests suggest can reach significantly higher accuracy than traditional, manual planning—in the 90% range versus the 70–80% range once considered a win. Mugsy will run this model alongside the existing inventory-planning process, using it first as a decision-support tool and, over time, as a forecasting engine, with people supervising and refining its recommendations.

CASE STUDY Mugsy

Looking ahead, Mugsy expects to move into the Run phase by turning some of those recommendations into semi-automated workflows—such as generating purchase-order suggestions based on current and projected demand, with human approval before anything is sent to suppliers. The longer-term vision for the Scale phase is to connect these decisions with marketing, product development, and other functions so that pricing, promotion, and inventory strategy can be adjusted in concert.

The Mugsy story illustrates a pattern that runs through the broader mid-market data: progress with AI starts when leaders frame the work around a business problem they care about—in this case, inventory, margin and clarity on performance. Beyond that, the new system should make it much easier to find, and comfortably invest in, new products and new sources of revenue.

> Ultimately, it comes down to unveiling opportunities that we either haven't thought of or just haven't had enough data to support investing in.
>
> – Leo Tropeano, Mugsy CEO –

Conclusion and Takeaways

Mid-market CEOs overwhelmingly believe AI can create value for their businesses, but most remain stuck in a prolonged experimentation phase, using pilots and point solutions rather than building AI into the core operating model. The survey data shows that while 98% of CEOs see AI's relevance and 60% already have active projects, only a small minority have a coordinated, company-wide AI strategy with clear ownership, shared data, and defined outcomes.

To close this gap between ambition and execution, CEOs will need to reframe AI from a collection of tactical tools to a strategic capability anchored in specific business results. That means shifting the primary question from "Should we use AI?" to "Which outcomes—margin, growth, working capital, customer lifetime value—should AI be responsible for improving?" In practice, this requires moving beyond an efficiency-only mindset, where most current initiatives focus on cost reduction and productivity gains, and instead connecting early wins to new products or sources of revenue and differentiated customer experiences.

The path forward lies in treating AI as part of how the business runs, not a set of side projects. Leaders can start by unifying critical data across functions to build a shared "corporate brain," then work through the Walk–Run–Scale sequence: first generating trustworthy, cross-functional intelligence, then deploying human-in-the-loop automation, and ultimately coordinating AI-driven decisions across sales, operations, finance and marketing. The Mugsy case study shows how starting with a specific outcome—freeing working capital and improving inventory performance—then mapping assets, building a data and semantic foundation, and phasing in AI-supported planning can turn abstract potential into measurable financial and operational gains.

For mid-market CEOs, the mandate is straightforward: stop adding more disconnected tools and invest instead in integration and governance that connects data, decision-making, and action. By grounding AI programs in problems the business actually cares about, building a unified intelligence layer, and gradually automating high-confidence decisions, leaders can turn today's scattered pilots into a virtuous cycle of learning and improvement, and convert AI from a tactical experiment into a durable source of competitive advantage.

VirtuousAI helps mid-market companies move beyond AI experimentation and transform how they operate with a strategic AI capability. Powered by its proprietary platform BAIO (Business Automation, Intelligence & Outcomes), VirtuousAI unifies data, decisions and actions into a continuous learning system that drives measurable outcomes and durable competitive advantage.

Learn how the connected intelligence of BAIO from VirtuousAI helps companies move faster, operate smarter, and scale effectively at **VirtuousAI.com**.

Chief Executive Group exists to improve the performance of U.S. CEOs, senior executives and public company directors, helping you grow your companies, build your communities and strengthen society. Through our peer networks, live events, proprietary research and flagship publications–including *Chief Executive* and *Corporate Board Member*–we connect CEOs, senior leaders and public company directors of nearly every sizable company in the United States, sharing your experiences, insights and hard-won wisdom to mutually benefit each other.

Learn more at **ChiefExecutiveGroup.com**.