The Agentic Workforce Economy

How Digital Labor Is Reshaping SMB Growth and Redefining the Role of IT Providers

Contents

INTRODUCTION

Agentic labor is an interconnected operating system that touches every dimension of the SMB business model at once.

For most of modern economic history, the growth of a small business followed a predictable logic: more revenue required more people. Headcount and output moved in rough proportion and the constraints of human labor (its cost, its scarcity, its geographic limits) defined the ceiling of what any small or medium-sized business could reasonably achieve. That relationship is now breaking. AI agents, automated workflows and intelligent software systems have introduced a new category of labor into the SMB economy; one that operates continuously, scales without hiring and performs cognitive tasks once reserved for human workers. The implications are structural, operational, financial and competitive, and they are already underway.

The macroeconomic signal is visible and accelerating. U.S. labor productivity has risen from 1.43 percent to 2.16 percent annually since late 2022, a pace not seen since the early years of the internet economy.[1] Deloitte's modeling tells us why: businesses that move from basic to intermediate AI adoption see profitability uplifts of roughly 45 percent. Those that reach full integration see uplifts of 111 percent.[2] The modeling is drawn from a study of 1,000 Australian SMBs; a context whose structural parallels to North American small business markets, including comparable firm-size distribution, digital infrastructure access and labor cost dynamics, make the directional findings applicable beyond their original geography. The curve is exponential.

Beneath the headline adoption figures lies a more complicated picture. Three quarters of SMBs are now investing in AI,[3] and more than half report active usage,[4] but adoption has proceeded unevenly and, in most cases, without the organizational foundations required for AI to deliver compounding returns. The divide is not between adopters and non-adopters. It runs between organizations that have deployed AI at the surface and those that have restructured their workflows, data infrastructure and governance models around it. That distinction, between using AI and operationalizing it, is the defining competitive variable of the coming decade.

Agentic labor, as this report defines it, is not a single technology or a discrete product category. It is an interconnected operating system that touches every dimension of the SMB business model at once. AI agents change hiring strategies, alter revenue-per-employee economics, create new demands on data infrastructure and require governance frameworks most small businesses haven't yet built. They also expand the cybersecurity attack surface in proportion to their value: 88 percent of SMB breaches now involve ransomware, far surpassing the 39 percent rate seen in larger enterprises.[5] The same agents that drive efficiency are the ones opening new vulnerabilities. Agentic labor is both the opportunity and the risk, and it demands a unified response rather than a piecemeal one.

1 Tsipursky, Gleb. AI Productivity Is Finally Hitting the Real Economy. The World Financial Review, 4 Mar. 2026.

2 Deloitte. The AI Edge for Small Business: Increased SMB AI Adoption Can Add $44 Billion to Australia's Economy. Deloitte Australia Press Room, 25 Nov. 2025.

3 Salesforce. AI and the Future of Small Business (A Trends Report Recap). Salesforce Blog, 2025.

4 McCabe, Laurie. How SMBs Are Adopting AI—And What Comes Next. SMB Group, 18 Aug. 2025.

5 Halcyon. Small and Medium Businesses Under Siege. Halcyon, 2025.

WHY NOW

Every major technology wave has a moment when the economics shift from experimental to operational, when the capability stops being impressive and starts being dependable enough to build a business around. That moment arrived for digital labor in the second half of 2025, and it arrived not because of a single breakthrough but because of four compounding inflections that happened in close succession.

The first was a capability threshold. Frontier models crossed into a qualitatively different range of performance on the tasks that matter most for business execution: coding, reasoning, multi-step tool use and sustained task completion across longer workflows. The relevant shift was operational reliability; models became good enough at turning intent into execution that the failure rate on real business tasks dropped below the threshold at which human oversight of every output was required. That change made autonomous action practical in a way it had not been before.

The second was the maturation of the agentic stack. Model capability alone does not produce deployable digital labor. What makes model gains usable in real operating environments is the surrounding infrastructure: orchestration layers, harnesses, execution frameworks, memory and context management, tool-calling protocols and the emerging standard of the Model Context Protocol (MCP) that allows agents to interact reliably with external systems. That stack matured rapidly alongside the models in 2025, and the combination; capable models running on deployable infrastructure; is what moved digital labor from demonstration to production.

The third was the emergence of the persistent agentic loop. The AI interactions most businesses have experienced to date are session-based: a human opens an interface, submits a prompt, receives a response, closes the session. That model is giving way to something structurally different. Always-on agents remain active between sessions. They monitor open commitments, surface next actions, request approvals, chase missing context and increasingly prompt the human rather than waiting to be prompted. That inversion; from AI-as-tool to AI-as-operating-layer; changes the rhythm of work itself, and it is already becoming visible in the workflows of the businesses that have moved furthest along the maturity curve.

The fourth was the arrival of agentic monitoring and operational resiliency as a practice category. Autonomous agents operating across business workflows require a control plane; observability into what agents are doing, governance over what they are permitted to do, recovery mechanisms for when they fail, and policy frameworks that define the boundaries of autonomous action. The emergence of that control plane is what makes managed intelligence a rinse-and-repeat service model rather than a bespoke deployment exercise. It is also what makes the managed intelligence provider an essential partner rather than an optional one: the businesses that deploy agents without the monitoring and resiliency layer are accumulating operational risk at the same rate they are accumulating productivity gains.

These four inflections are not independent trends. They compound. The capability threshold made more work executable. The agentic stack made it deployable. The persistent agentic loop changed how work is initiated and coordinated. Monitoring and resiliency make it governable and repeatable at scale. Together, they mark the beginning of a phase change in the SMB economy, the conditions under which AI adoption becomes workforce re-architecture. The rest of this report examines what that re-architecture looks like, why most SMBs are not yet capturing it and what it means for the technology partners who serve them.

MSP confidence in their ability to guide SMBs on AI deployment has fallen from 90 percent to roughly 50 percent in a single year, even as demand has surged.

At the center of this transformation is the technology partner ecosystem. SMBs are not, and largely cannot be, the primary architects of their own agentic labor strategies. AI expertise has become the third most important attribute SMBs seek in a managed service provider (MSP), behind only threat prevention and round-the-clock support, yet fewer than half of MSPs have built or deployed AI-specific capabilities for their clients.[6] MSP confidence in their ability to guide SMBs on AI deployment has fallen from 90 percent to roughly 50 percent in a single year, even as demand has surged.[7] The result is a structural supply-demand mismatch at precisely the moment the stakes are highest.

That gap is giving rise to a new category of provider. As SMBs deploy digital workers, the managed service providers who serve them are being asked, and in many cases compelled, to evolve into what Pax8 has termed Managed Intelligence Providers (MIPs): technology partners who do more than merely manage infrastructure but instead, orchestrate the intelligence flows that run the business. The MIP is not an upgraded MSP. It is a fundamentally different role; closer to a fractional COO for the AI-native SMB than to a traditional technology vendor.

Where the MSP managed systems, the MIP manages outcomes: deploying and optimizing AI agents, governing autonomous decision systems, integrating data environments, and ensuring that the operational logic underlying the agentic workforce is performing, compliant, and compounding in the right direction. The economics reflect the stakes. AI services in the managed services sector are growing at 59 percent annually, compared to 13 percent for traditional managed services.[8]

The margin structure, the value proposition and the competitive dynamics of the channel are all being rewritten at the pace of the clients who are already running digital workforces that nobody is governing.

Pax8's own Q1 2026 SMB Technology Pulse, surveying 400 U.S. small business leaders involved in technology decisions, finds that 84 percent would trust an outside technology advisor to guide their AI implementation, and 70 percent agree that outside partnerships are necessary to fully benefit from AI.[9] The demand for managed intelligence guidance has arrived, and SMBs are actively looking for partners capable of meeting it.

Drawing on global market research, proprietary Pax8 ecosystem data and original analysis, this report maps the forces reshaping how small businesses grow, compete and operate in an economy where digital labor is no longer a future scenario. It covers the productivity and valuation dynamics unlocked by AI and automation; the perception gaps, readiness deficits and governance failures that hold most SMBs back; the cybersecurity risks that scale alongside every new deployment; and the evolving role of technology partners as the infrastructure through which SMB AI success is built and sustained.

The central thesis is straightforward, even if its implications are not: agentic labor is not a future scenario to prepare for. It is a present restructuring to navigate. The SMBs and technology providers that understand its architecture; that recognize digital labor as a system requiring data foundations, security governance, workflow redesign and intelligent partnership; will define the next operating model for small business growth. The race to managed intelligence is not a matter of technology readiness alone. It is a matter of organizational will, strategic clarity and the quality of the partnerships that small businesses are willing to build. The businesses and technology providers that grasp this distinction first will define what the SMB economy looks like for decades to come.

84%

Would trust an outside advisor to implement AI.

62%

Say AI required for competitiveness within three years.

70%

Agree outside partners are necessary to fully benefit.

74%

Believe AI lets them compete with larger companies.

6 OpenText. OpenText Cybersecurity Finds 92% of Managed Service Providers See AI-Driven Growth, but Readiness Gap Widens. OpenText Press Releases, 24 Sept. 2025.

7 OpenText. OpenText Cybersecurity Finds 92% of Managed Service Providers See AI-Driven Growth, but Readiness Gap Widens. OpenText Press Releases, 24 Sept. 2025.

8 McBain, Jay. MSPs in the AI Era: Your Guide to Capturing New Opportunities. Channel Insights, SoundCloud, 2025.

9 Pax8. "SMB Technology Pulse Survey: Q1 2026 Topline Summary." Pax8, Q1 2026. Proprietary research. Data on file.

Part 01.

The New Workforce Economy

Chapter 01.

Headcount-Free Growth: The New Valuation Metric

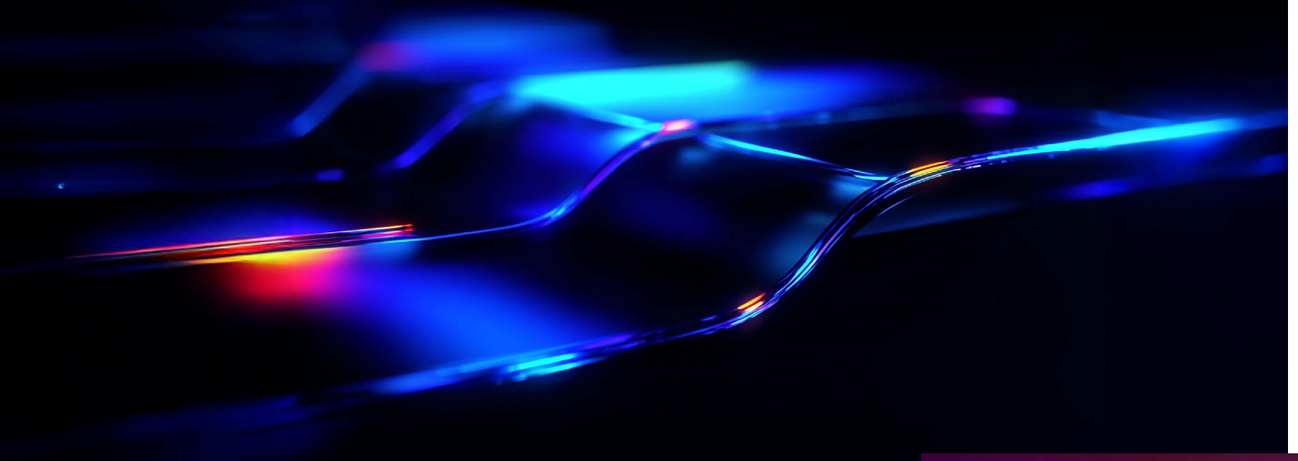

The productivity gains from digital labor are well documented. The valuation implications are not.

When a small business grows revenue without growing headcount proportionally, something changes beyond its cost structure. It changes the story that business tells every outside observer evaluating it: lenders assessing creditworthiness, acquirers modeling multiples, investors weighing risk. According to IDC's *Worldwide Small and Medium-Sized Business 2025 Predictions*, SMBs that successfully demonstrate the ability to break the direct correlation between revenue growth and headcount growth will achieve higher valuations and attract greater outside investment; a structural shift that reframes digital labor from an operational tool into the foundation of a new growth model.[10]

This is a structural assertion about what outside capital pays for, and it reframes digital labor from a productivity tool into something closer to a financial strategy. Revenue-per-employee is the ratio at the center of this shift. Most SMB owners track it, if at all, as an HR metric: a gauge of workforce efficiency. It is becoming something more consequential: a signal to lenders, acquirers and equity partners that the business has found a way to scale output without the friction of proportional hiring.

When that ratio rises consistently, it tells a story that commands different terms. This chapter examines what drives that ratio, who is moving it and what the macroeconomic data already reveals about the gap between the businesses capturing these gains and the majority that are not.

THE ADOPTION VELOCITY NOBODY IS TALKING ABOUT

The U.S. Census Bureau's Business Trends and Outlook Survey, which tracks AI adoption biweekly across approximately 1.2 million U.S. businesses, documents the transition in real time,[11] with businesses using AI across any business function reaching 17.3 percent by November 2025.[12] When a company can grow revenue without growing headcount proportionally, it changes two things simultaneously: the cost structure of scaling, and the story that business tells to any outside observer evaluating it. Both matter enormously for small businesses that have historically been locked out of the valuation multiples and investment interest available to enterprises with demonstrated scalability.

REVENUE-PER-EMPLOYEE AS A CAPITAL STRATEGY

That ratio, revenue-per-employee, is the mechanism through which headcount-free growth becomes a capital story. When it rises consistently, it signals to any external evaluator that the business has found a way to scale without the friction of proportional staffing. For an SMB owner considering a sale, a refinancing or an equity raise, that signal is worth real money. For a technology partner helping that SMB deploy digital labor, it represents the language through which to frame the value of the work. It measures how much output a business generates per unit of human labor cost, and when it rises consistently, it signals to any external evaluator that the business has found a way to scale without the friction of proportional staffing.

[this is] where the right technology partner delivers the most differentiated value.

Deloitte's modeling of AI maturity levels among SMBs puts the profitability dimension in sharp relief. Businesses that move from basic to intermediate AI adoption see profitability uplifts of approximately 45 percent.[13] Those that progress from intermediate to fully enabled integration see uplifts of approximately 111 percent.[14] The steepest part of this exponential return lies in the second half of the journey. This is precisely the segment where most SMBs have not yet arrived, and where the right technology partner delivers the most differentiated value.

What connects the profitability data to the valuation thesis is the mechanism: agentic labor absorbs output without absorbing headcount. When an AI agent handles invoice reconciliation, customer follow-up, appointment scheduling or first-line IT support, those functions happen at near-zero marginal cost. Revenue grows. Payroll does not. The ratio moves. The business becomes worth more, not metaphorically, but in the real numbers that govern small business financing and acquisition.

10 Blackwell, Jason, et al. "IDC FutureScape: Worldwide Small and Medium-Sized Business 2025 Predictions." IDC, October 2024.

11 U.S. Census Bureau."Business Trends and Outlook Survey (BTOS)."

12 Goldschlag, Nathan. "How Many Businesses Are Using AI?" Agglomerations, 19 Dec. 2025.

13 O'Mahony, John, et al. "The AI Edge for Small Business: Increased SMB AI Adoption Can Add $44 Billion to Australia's Economy." Deloitte Access Economics, 25 Nov. 2025.

14 O'Mahony, John, et al. "The AI Edge for Small Business: Increased SMB AI Adoption Can Add $44 Billion to Australia's Economy." Deloitte Access Economics, 25 Nov. 2025.

THE MACROECONOMIC SIGNAL AND THE SMB GAP

The aggregate data documents the opportunity; it also documents who is not yet capturing it. Research finds that industries with higher AI-driven time savings are seeing real ROI, with 2.7 percentage points faster productivity growth per point of time saved.[15] Yet, only 10 percent of SMBs have fully integrated digital tools across their operations, despite 92 percent reporting some level of digital tool usage.[16] The vast majority of small businesses are in the messy middle: they have spent money on technology, but they have not unlocked the compounding returns that come from full integration.

McKinsey's 2025 State of AI report adds another layer of precision to this picture. While 88 percent of organizations now use AI in at least one business function, only 39 percent report EBIT impact at the enterprise level.[17] The gap between using AI and capturing enterprise-level financial returns from it is not a technology gap. McKinsey's high-performer analysis identifies the distinguishing factors as workflow redesign, talent strategy, data infrastructure and process integration, organizational decisions, not technical ones. This finding matters for SMBs because it explains why tool adoption alone does not produce the valuation outcomes that full integration does. The "last mile" of AI transformation, where technical capability must meet organizational redesign, is where value is captured or lost.

The competitive stakes are not lost on small business leaders themselves. Pax8's Q1 2026 SMB Technology Pulse finds that 62 percent of SMBs agree that without AI their business will not remain competitive within three years, and 74 percent believe AI gives small businesses the ability to compete directly with larger companies.[18] That conviction is running ahead of the organizational capability to act on it. Sixty-eight percent express at least moderate concern about falling behind competitors due to technology gaps, even as 97 percent report that technology spending has increased or held steady over the past year.[19] The ambition and the investment are present. The architecture that converts both into compounding returns is not.

Among SMBs currently using AI, 86 to 90 percent report that it makes operations more efficient and helps them scale services.[20] But 95 percent report needing more training to use AI effectively. The value is understood. The capability to capture it is not yet distributed. That gap between recognized opportunity and realized outcome is the specific space where a Managed Intelligence Provider operates.

THE MESSY MIDDLE: SMB AI ADOPTION SPECTRUM

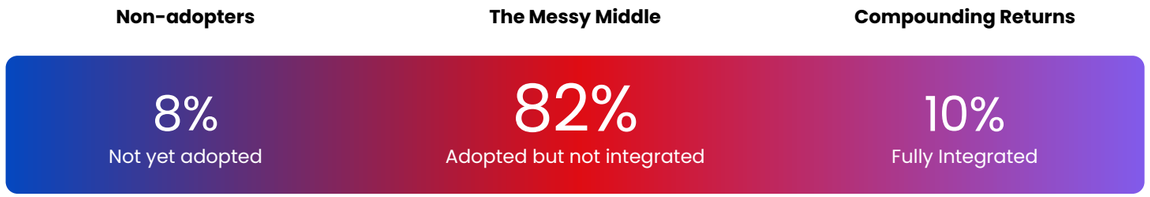

15 Bick, Alexander, Adam Blandin, and David Deming. "The State of Generative AI Adoption in 2025." Federal Reserve Bank of St. Louis, 19 Nov. 2025.

16 Alegbeh, Alchad, and Marvin Cruz. "Digital Transformation: How Small Businesses in Canada Are Leveraging AI and Technology for Growth and Productivity." Canadian Federation of Independent Business, 29 Sept. 2025.

17 McKinsey & Company. "The State of AI in 2025." McKinsey Global Institute, Nov. 2025.

18 Pax8. "SMB Technology Pulse Survey: Q1 2026 Topline Summary." Pax8, Q1 2026. Proprietary research. Data on file.

19 Pax8. "SMB Technology Pulse Survey: Q1 2026 Topline Summary." Pax8, Q1 2026. Proprietary research. Data on file.

20 Salesforce. "Small and Medium Business Trends Report, 6th Edition." Salesforce, Dec. 2024.

21 Railton, Matt. "IDC Defines Digital Labor for the Agentic Enterprise." No Jitter, 15 Oct. 2025.

22 Railton, Matt. "IDC Defines Digital Labor for the Agentic Enterprise." No Jitter, 15 Oct. 2025.

23 U.S. Small Business Administration, Office of Advocacy. "Frequently Asked Questions About Small Business." SBA, Feb. 2026.

World Bank. "Small and Medium Enterprises (SMEs) Finance." World Bank Group, 2025.

WHAT HEADCOUNT-FREE GROWTH ACTUALLY LOOKS LIKE

The practical expression of headcount-free growth is not dramatic disruption. It begins in the back office. IDC's research on digital labor finds that the roles being absorbed earliest are not the complex, judgment-intensive work that defines most small businesses' competitive differentiation, but rather, the operational functions that consume time without generating proportional value, such as: reconciling accounts, scheduling and rescheduling, triaging incoming requests, generating routine reports, following up on outstanding payments.[21] These are the functions that consume hours, require consistency and benefit from scale; precisely where software agents outperform humans on a cost-per-unit basis.

> 95 percent [of SMBs] report needing more training to use AI effectively.

Additional workforce data from IDC captures the consequence of this shift at the organizational level: in businesses that have invested in digital labor, 23 percent of roles have been newly created and 24 percent have been reassigned.[22] Nearly half the workforce in these organizations has been restructured around AI, not reduced by it. The narrative of agentic labor as pure headcount elimination misses the more important story: it is workforce reconfiguration at a pace that most businesses cannot manage without external guidance.

THE STRATEGIC IMPLICATION FOR TECHNOLOGY PARTNERS

The shift from headcount-based growth to output-based growth changes what SMBs need from their technology partners, and how those partners should frame the conversation. The traditional MSP value proposition was built around infrastructure reliability, security and cost management. These remain essential. But they do not speak the language of the business owner who is trying to grow revenue without a proportional payroll increase, or who is thinking about what their business will be worth in five years. The technology partner who reframes its value proposition around growth-without-headcount architecture, positioning agentic labor deployments not as cost savings but as valuation multipliers, is speaking directly to what this class of SMB is optimizing for:

- Build reporting frameworks that track revenue-per-employee alongside traditional IT metrics.
- Help clients document their headcount-free growth story in terms that translate to lenders and acquirers.
- Position each digital labor deployment not as a software implementation but as a contribution to a capital strategy.

The SMB that can demonstrate to a banker or a buyer that its revenue grew 30 percent while headcount grew 5 percent is telling a fundamentally different story about its business, and one that commands fundamentally different terms.

THE FLOOR HAS MOVED

The historic correlation between revenue growth and headcount growth is not breaking gradually. It is breaking now, at a measurable pace, in data that is federally and globally tracked and publicly available. The businesses at the leading edge of this transition are the early visible portion of a wave that the broader SMB economy is beginning to enter.

> Software agents outperform humans on a cost-per-unit basis.

The constraint on how widely and how quickly the SMB economy captures these gains is not the availability of the technology. Globally, the World Bank estimates that small and medium-sized enterprises represent approximately 90 percent of all businesses and more than 50 percent of employment worldwide.[23] The overwhelming majority of these businesses are not yet on the steep part of the AI value curve. The constraint is the organizational work required to get there: the workflow redesign, the data preparation, the governance structure, the ongoing optimization and the availability of partners capable of doing that work at scale. But before any of that work can begin, there is a more fundamental obstacle to clear. It's not budget, access or the technology itself.

Chapter 2 examines the widespread belief that this transformation does not apply to small businesses like theirs.

Chapter 02.

The Applicability Illusion: Why SMBs Think This Doesn't Apply to Them

There are more than 400 million small businesses in the world.[24] The overwhelming majority of them are not sitting on the sidelines of the AI economy because they lack the budget. They are not waiting because they lack access to the tools. They are waiting because they do not believe the transformation applies to them: to a two-person landscaping company, a four-chair hair salon, a family-owned HVAC distributor in a mid-sized city. The technology, in their view, is for someone else's kind of business.

The data confirms it, and it is consistent across borders. In fact, the single most cited barrier to generative AI adoption among non-adopters was unsuitability to the SMB's work, cited by 57 percent of businesses that have not adopted it.[25]

What makes this perception so consequential is not that it is held; it is that it is wrong. In precisely the same size categories, in precisely the same sectors, businesses that have crossed the applicability threshold are reporting 29 percent average productivity gains in the first year of digital tool adoption and a return of $1.60 for every dollar invested.[26] The businesses achieving these returns are not structurally different from those sitting on the sidelines. They are demographically similar, operationally comparable and often competitors in the same local markets. The variable separating them ultimately comes down to conviction.

This chapter examines that gap. Not the technology gap, which is closing rapidly and unevenly, but the perception gap that precedes it: the belief, widely held and empirically indefensible, that the AI transformation reshaping the global economy does not reach the businesses that make up 90 percent of it. This is the applicability illusion. And it is the largest single obstacle standing between most small businesses and the value curve examined in Chapter 1.

"Businesses that have crossed the applicability threshold are reporting 29 percent average productivity gains in the first year."

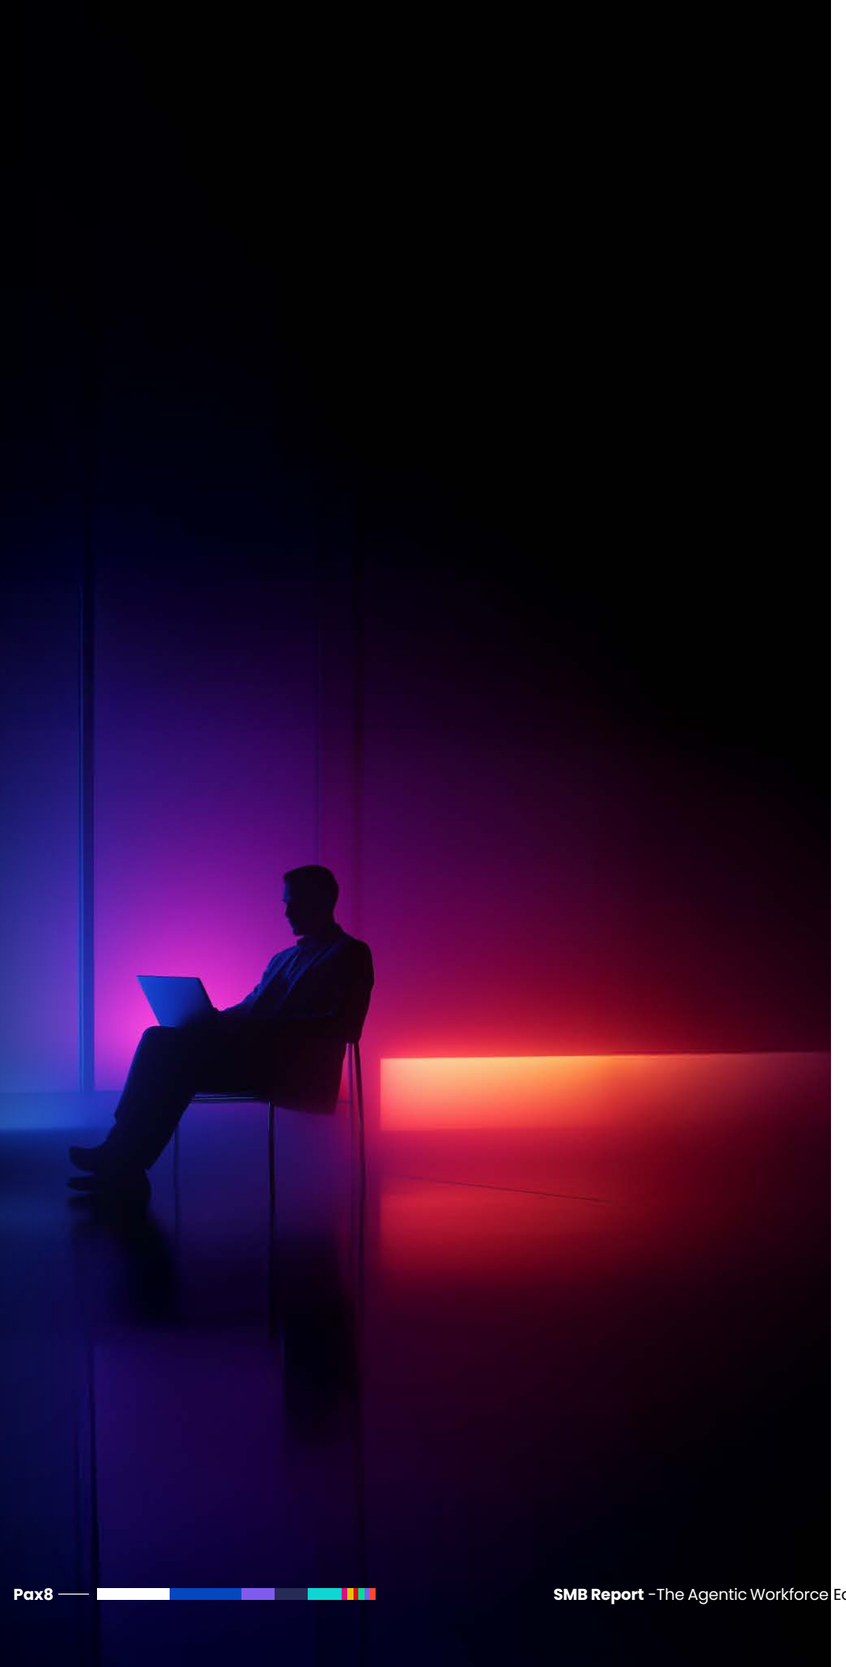

THE PERCEPTION GAP, QUANTIFIED

The applicability illusion is not a fringe position. It is the dominant view among the smallest segment of the SMB economy: the micro-businesses that constitute the vast majority of firms by count. A recent survey of more than 5,000 SMBs across seven countries found that the single most cited barrier to generative AI adoption among non-adopters was unsuitability to the work, not cost.[27] A categorical judgment that the technology had no relevance to their specific business context. The finding held across industries and geographies; a perception gap that is not a quirk of any single market but a structural feature of the global SMB economy.

Among small businesses that have not yet adopted AI, nearly three quarters say they would move forward if presented with clearer evidence of return on investment, not a general capability pitch, but proof that the technology delivers in an environment like theirs.[28] This is the logical response of a business owner who has never seen AI work in an environment like theirs, served by an industry that has often led with capability rather than context. The burden of proof has shifted to the provider ecosystem, and most of that ecosystem has not yet met it.

The consequences of this demand-supply mismatch at scale bears this claim out. Among SMBs not currently using AI, the leading barriers are not financial: they include lack of relevance to their business, limited internal skills and security and privacy concerns.[29] The perception of irrelevance leads. Budget follows. This ordering matters because it determines the nature of the intervention required. A cost problem is solved with pricing. A skills problem is solved with training. A perception problem is solved with evidence, and evidence, in this context, means stories from businesses that look like theirs.

The most striking evidence of the perception gap is in the divergence between how AI users and non-users see the same landscape. Among small businesses currently using AI, nearly four in five describe the technology as essential to their competitiveness; a conviction shared by fewer than half of those who have not yet adopted it.[30] The same market, viewed through two different lenses, produces almost incompatible readings of reality. This is a difference in experiential reference points, one that compounds over time as the gap between adopters and non-adopters widens.

24 Kumar, Naveen. "Small Business Statistics of 2026 (New Trends & Data)." DemandSage, 1 Jan. 2026.

25 OECD. "Generative AI and the SME Workforce: New Survey Evidence." OECD Publishing, Paris, 2025.

26 Alegbeh, Alchad, and Marvin Cruz. "Digital Transformation: How Small Businesses in Canada Are Leveraging AI and Technology for Growth and Productivity." Canadian Federation of Independent Business, 29 Sept. 2025.

27 OECD. "Generative AI and the SME Workforce: New Survey Evidence." OECD Publishing, Paris, 2025.

28 Reimagine Main Street / Public Private Strategies Institute. "Beyond Efficiency: Small Businesses Look to AI for Competitive Edge." PayPal, June 2025.

29 McCabe, Laurie, and Sanjeev Aggarwal. "The Impact of AI on SMBs: 2025 Trends, Challenges, and Opportunities." SMB Group, June 2025. Sponsored by Workday.

30 Reimagine Main Street / Public Private Strategies Institute. "Beyond Efficiency: Small Businesses Look to AI for Competitive Edge." PayPal, June 2025.

WHAT ADOPTION ACTUALLY LOOKS LIKE AT THE SMALLEST SCALE

Part of the applicability illusion is sustained by a mismatch between the mental image of AI adoption and its actual practice among small businesses. The dominant cultural image of AI deployment involves enterprise infrastructure, data science teams and multi-year implementation programs. The actual experience of an SMB that has crossed the adoption threshold looks almost nothing like this.

The average SMB using generative AI reported saving 1.08 hours per day, more than an hour of reclaimed capacity.

The pattern holds globally. A recent analysis of 160 million professionals across more than 18 million small businesses found that the top generative AI use cases among adopters were not operational overhauls but incremental efficiency gains: automating repetitive tasks, simplifying processes, writing reports and memos and data entry; the same back-office time recaptures documented at the national level. The Canadian Federation of Independent Business, surveying 1,683 SMB owners, put a precise figure on the return: the average SMB using generative AI reported saving 1.08 hours per day, more than an hour of reclaimed capacity for every working day, redirected toward higher-value activities.[31] For a sole proprietor who currently spends three hours a day on administrative tasks that require no judgment, this is a structural change in what their working day looks like.

The highest-rated AI benefits among current users are improved data analysis and decision-making (43 percent), better planning and forecasting capability (38 percent), and the ability to summarize information from multiple sources and eliminate repetitive tasks (31 percent each).[32] These compose the daily operational frictions that consume disproportionate time in every small business, regardless of industry. The landscaper who spends ninety minutes each morning responding to estimate requests, the salon owner who manually reconciles appointment no-shows against payroll, the HVAC distributor who compiles weekly inventory reports by hand: these are the lived experiences of small business ownership, and they are precisely the functions where digital labor delivers its clearest early returns.

The adoption gap between large and small businesses, which once measured in years, has compressed dramatically. The Small Business Administration's (SBA) longitudinal analysis of Business Trends and Outlooks Survey (BTOS) data found that in early 2024, large businesses used AI at 1.8 times the rate of small businesses.[33] By August 2025, the gap had narrowed to approximately 1.2 times.[34] Previous technology adoption cycles; broadband internet, cloud computing, mobile payments; saw small businesses lag their larger counterparts by years. The current compression is without precedent in the history of SMB technology adoption, and it is driven primarily by the availability of zero-infrastructure, zero-IT-staff tools that deliver immediate, visible utility.

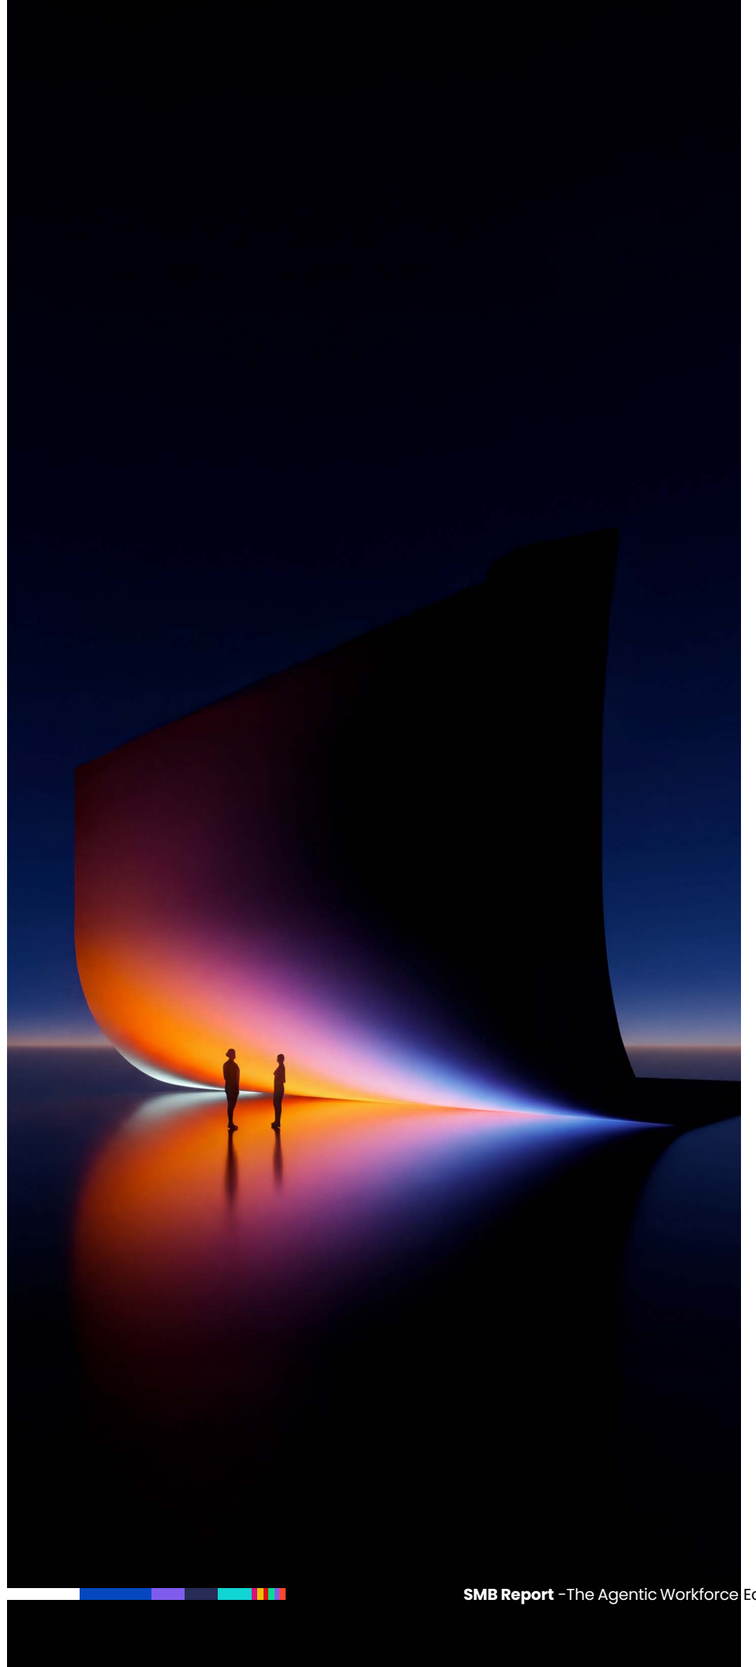

THE $1.60 RETURN NOBODY IS TALKING ABOUT

As mentioned earlier, there is a predicted $1.60 return for every dollar invested in digital tools across the SMB segment, with a 29 percent average productivity gain in the first year of adoption.[35] Those businesses that have fully integrated tools across all core functions, report $2.40 in returns for every dollar invested, or 1.7 times the return of less mature adopters. And these returns are not confined to any particular industry or business type, spanning professional services, retail, food service, construction, healthcare and personal services; the full range of sectors that comprise the micro and small business economy.

The 55 percent of businesses that saw returns on their technology investment within the first two years did not do so by deploying sophisticated enterprise AI infrastructure.[36] They did so primarily by automating the subset of their daily operations that consumed hours without generating proportional value; the same back-office functions identified in Chapter 1 as the earliest absorption targets for digital labor. The path to the $1.60 return begins with a decision that requires no technical expertise, no capital infrastructure and no IT staff. It requires the conviction that the tools have something to offer, applied to a process that consumes time and delivers measurable output.

The AI-use data from the SMB Group survey reinforces the same pattern from a different angle. Among SMBs reporting high or transformational AI value, the functions generating the clearest returns are IT management, finance and HR: the administrative and operational core of any business, regardless of what that business does or sells.[37] The applicability illusion draws its power from a category error: the belief that because AI does not change what a business makes or sells, it has nothing to offer. The evidence consistently shows that the value lives not in the product or service but in the operations surrounding it, and those operations are, with few exceptions, identical across industries.

31 LinkedIn Economic Graph. "How Small Businesses Can Win in 2026: Work Change Special Report." LinkedIn, Dec. 2025.

Alegbeh, Alchad, and Marvin Cruz. "Digital Transformation: How Small Businesses in Canada Are Leveraging AI and Technology for Growth and Productivity." CFIB, 29 Sept. 2025.

32 Alegbeh, Alchad, and Marvin Cruz. "Digital Transformation: How Small Businesses in Canada Are Leveraging AI and Technology for Growth and Productivity." Canadian Federation of Independent Business, 29 Sept. 2025.

33 Alegbeh, Alchad, and Marvin Cruz. "Digital Transformation: How Small Businesses in Canada Are Leveraging AI and Technology for Growth and Productivity." Canadian Federation of Independent Business, 29 Sept. 2025.

34 U.S. Chamber of Commerce. "Empowering Small Business: The Impact of Technology on U.S. Small Business." U.S. Chamber of Commerce, 13 Aug. 2025.

35 Alegbeh, Alchad, and Marvin Cruz. "Digital Transformation: How Small Businesses in Canada Are Leveraging AI and Technology for Growth and Productivity." Canadian Federation of Independent Business, 29 Sept. 2025.

36 Alegbeh, Alchad, and Marvin Cruz. "Digital Transformation: How Small Businesses in Canada Are Leveraging AI and Technology for Growth and Productivity." Canadian Federation of Independent Business, 29 Sept. 2025.

37 McCabe, Laurie, and Sanjeev Aggarwal. "The Impact of AI on SMBs: 2025 Trends, Challenges, and Opportunities." SMB Group, June 2025. Sponsored by Workday.

The most common actions SMBs are currently taking to close their AI skills gap are telling: 68 percent are learning more about AI features in applications they already use.

AN EDUCATION CRISIS IN TECHNOLOGY CLOTHING

The consistent finding across every dataset examined in this chapter is that the primary barrier to SMB AI adoption is neither financial nor technical, but a knowledge gap. Businesses do not know what AI looks like when it works in an environment like theirs. They have not seen it. They have not heard a story that made it feel real, proximate and relevant to the decisions they make on a Tuesday morning.

To underscore this point, among businesses with three to 19 employees, the segment where the applicability illusion is most concentrated, only eight percent report significant AI proficiency, and 39 percent report only internal experimentation with no external application.[38] Thirteen percent report no AI proficiency at all. These are not businesses that have evaluated AI and found it wanting. They are businesses that have not yet found a reason to look.

Pax8's own discovery research, conducted with 18 small and medium business decision-makers across three segments (those running IT in-house and intending to stay that way, those running IT in-house and looking to partner with a managed service provider and those already tethered to one) reinforces the applicability illusion as a present and consistent feature of the SMB economy. Across all three segments, AI adoption sat at roughly the same maturity level: occasional use of consumer-grade tools for drafting, summarization and basic research, without a formal strategy, deliberate integration into operating workflows or any organizational sense that the technology was structurally relevant to the business. The interviews surfaced the same explanation that the global survey data documents at scale. These were not businesses that had evaluated agentic labor and found it lacking. They were businesses that had not yet encountered AI working in an environment that resembled theirs, and the absence of contextual proof was enough to keep the conviction gap intact even among the smallest businesses where the case for absorption is, on paper, the most direct.[39]

The most common actions SMBs are currently taking to close their AI skills gap are telling: 68 percent are learning more about AI features in applications they already use, and 63 percent are assessing where AI can create the most value in their specific context.[40] These are orientation activities, not deployment activities. The market is in a sustained phase of trying to understand before it commits; a phase that has persisted despite two years of mainstream AI availability because the understanding required is contextual, not general. A small business owner does not need to understand large language models. They need to understand what one looks like running their invoicing workflow or their customer follow-up queue. That contextual translation has been the missing piece.

The data also captures the agentic AI dimension of this knowledge gap in particularly sharp terms. While 90 percent of SMB decision-makers report being either very or somewhat familiar with general AI, only 30 percent say the same about agentic AI: the autonomous, action-taking form that represents the next wave of digital labor deployment.[41] The gap between general AI familiarity and agentic AI familiarity is the gap between knowing that this technology exists and understanding what it can actually do in an operational context. As agentic labor shifts from assistive to autonomous, that gap will determine which businesses are positioned to capture the compounding returns of full integration and which remain stranded in the experimentation phase.

What makes this knowledge gap particularly durable is that the hardest parts of AI adoption are not the models, but the organizational conditions surrounding them. IDC's analysis of enterprise AI adoption identifies skills gaps, integration into legacy workflows, compliance uncertainty and governance as the dominant obstacles even among organizations that have made significant technology investments.[42] For SMBs, which have fewer internal resources, less formal governance infrastructure and thinner technical bench strength than enterprise organizations, those frictions are likely to be sharper rather than more forgiving. The implication is significant: closing the applicability illusion does not end the adoption challenge; it reveals the next one. Once an SMB owner accepts that agentic labor applies to their business, they immediately encounter the implementation frictions that require external expertise to navigate. The education gap and the execution gap are sequential problems, and technology partners who can address both in sequence are positioned to own the relationship through the full adoption journey.

38 McCabe, Laurie, and Sanjeev Aggarwal. "The Impact of AI on SMBs: 2025 Trends, Challenges, and Opportunities." SMB Group, June 2025. Sponsored by Workday.

39 Galvan, Moriah. "SMB Discovery." Pax8 UX Research, May 2026. Proprietary research. Data on file.

40 McCabe, Laurie, and Sanjeev Aggarwal. "The Impact of AI on SMBs: 2025 Trends, Challenges, and Opportunities." SMB Group, June 2025. Sponsored by Workday.

41 McCabe, Laurie, and Sanjeev Aggarwal. "The Impact of AI on SMBs: 2025 Trends, Challenges, and Opportunities." SMB Group, June 2025. Sponsored by Workday.

42 Schubmehl, David, and Lange, Kathy. "What Every Company Can Learn from Frontier Firms Leading the AI Revolution: Accelerating Innovation with AI". IDC, Nov. 2025. Sponsored by Microsoft.

THE LEAPFROG OPPORTUNITY HIDDEN IN THE APPLICABILITY ILLUSION

The applicability illusion is predominantly a problem of the smallest organizations. But among those organizations, there is a paradox that the data is beginning to surface: the smallest businesses may be the most structurally positioned to capture AI's gains, precisely because they have the least to unlearn.

Enterprise AI transformation is expensive and slow not primarily because the technology is complex, but because deploying it requires dismantling decades of entrenched process, legacy infrastructure and organizational habit. Research has found that the primary obstacle to enterprise-level AI value was not model quality or data availability but the organizational redesign required to convert localized experiments into standard operating models.[43] For a fifty-person professional services firm with no ERP system, no legacy CRM architecture and no change management bureaucracy, that obstacle does not exist. The absence of infrastructure that larger businesses must work around is, for the smallest SMBs, a structural advantage.

IDC's 2026 SMB digital landscape research identifies this dynamic as a defining feature of the current adoption wave. Micro-businesses that build their operational infrastructure on AI-first platforms from the outset; bypassing the ERP-to-cloud migration journey that consumes years and capital for mid-market firms; are achieving a level of automation that would have been inaccessible to businesses of their size in any prior technology cycle.[44] The traditional adoption curve, in which larger firms lead and smaller ones follow years later, is being compressed and in some segments inverted. The born-automated small business; one that has never known a different way of operating; is an emerging category, and it will compound its advantage with every passing quarter.

The implication for technology partners is significant. The most valuable long-term client relationships in the SMB market may not be the established mid-market organization upgrading legacy systems. They may be the nascent micro-business that has not yet built its operational infrastructure; the one that will build it once, correctly, around agentic labor from the outset, and will never need to rebuild. Capturing that client at formation is a compounding investment in the highest-growth segment of the SMB economy.

BREAKING THE ILLUSION: WHAT WORKS IN PRACTICE

If the applicability illusion is fundamentally an education problem, then the interventions that break it are fundamentally educational ones; not in the classroom sense, but in the sense of making an abstract possibility feel concrete, specific and achievable. The data on what moves businesses from non-adoption to adoption points consistently toward three mechanisms.

The first is peer evidence. SMBs that have seen businesses identical to theirs (same size, same sector, same operational profile) achieve measurable returns from AI adoption are significantly more likely to invest than those who have only encountered general market claims about AI's potential. This is the logic behind industry-specific case studies and the reason generic capability demonstrations consistently fail to move the needle with non-adopters.

The second is low-barrier entry. The research identifies the most effective on-ramps to AI adoption as embedded features in existing applications: meeting summaries, content suggestions, automated follow-up sequences, rather than standalone AI deployments.[45] The business owner who discovers that the accounting software they already pay for will now draft their collection emails has encountered AI in its most persuasive form: a tangible time saving delivered inside a familiar context, with no new vendor relationship required.

The third is contextualized ROI. The $1.60 return and the 29 percent productivity gain are not abstractions to a business owner who has been shown, specifically, which hour of their day those numbers correspond to. The applicability illusion dissolves when they are shown what thirty fewer minutes of daily administrative work looks like in annual terms: fifteen days of reclaimed capacity per year, at whatever their effective hourly rate translates to.

For technology partners, especially MIPs, this translates into a specific practice implication. The most effective sales motion for the non-adopter market is not product demonstration, but mirror-selling: showing a prospect, specifically, what a business that looks exactly like theirs is doing with agentic labor and what it has produced. The 82 percent of smallest-firm owners who believe AI does not apply to them represent the largest untapped addressable market in the SMB technology economy. They are not a lost cause. They are an education problem awaiting a credible teacher.

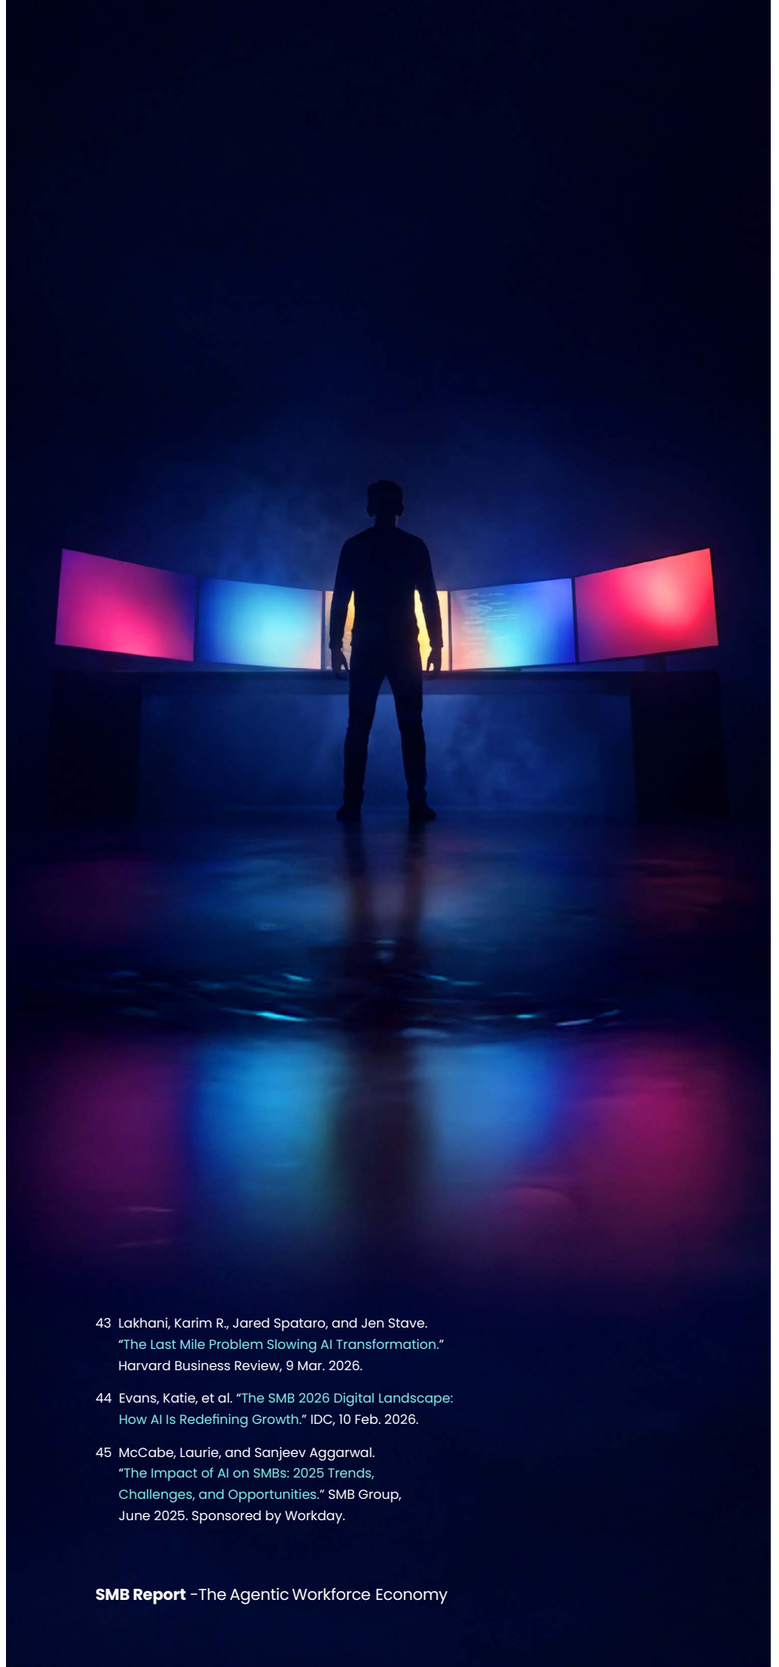

43 Lakhani, Karim R., Jared Spataro, and Jen Stave. "The Last Mile Problem Slowing AI Transformation." Harvard Business Review, 9 Mar. 2026.

44 Evans, Katie, et al. "The SMB 2026 Digital Landscape: How AI Is Redefining Growth." IDC, 10 Feb. 2026.

45 McCabe, Laurie, and Sanjeev Aggarwal. "The Impact of AI on SMBs: 2025 Trends, Challenges, and Opportunities." SMB Group, June 2025. Sponsored by Workday.

THE COST OF WAITING

Every quarter that a non-adopting SMB waits, the businesses around it that have crossed the threshold are widening their operational advantage: serving more clients with the same headcount, recapturing hours that compound into capacity, building the data foundations and workflow habits that accelerate further adoption. The gap between the 10 percent that have fully integrated digital tools and the 90 percent that have not is measured in operating model maturity, and it is growing.

For most small businesses in any market, the risk is discovering too late.

For most small businesses in any market, the risk is discovering too late that the competitor across the street built their operations around digital labor while they were still deciding whether it applied to them. The businesses that will feel this most acutely are not the early adopters, who are already on the compounding part of the curve, or the determined holdouts, who have made a considered choice. They are the businesses in the messy middle: those who have heard enough to be curious, invested enough to have some tools running, but not yet committed enough to have restructured anything fundamental around them.

Chapter 3 examines the macroeconomic evidence for why the messy middle is the most dangerous place to be, and why the industries documenting the fastest productivity growth are precisely those where AI-driven time savings have moved from experimentation to standard operating practice.

Chapter 03.

The Macro Signal Is Already Here. SMBs Are Missing It.

The productivity argument for agentic labor is no longer theoretical. It is federally documented, measurable at the industry level and large enough to have shifted the national economic data. More small businesses than ever are using artificial intelligence. Vanishingly few have operationalized it. The distance between those two statements is the distance between the businesses that are capturing the productivity surge and the majority that are watching it happen to someone else.

Generative AI reached approximately 53 percent population-level adoption within three years of its mass-market introduction; faster than the personal computer or the internet.[46] The speed of that diffusion is without modern precedent, and its economic consequences are already visible: U.S. consumer surplus from generative AI reached an estimated $172 billion annually by early 2026, up from $112 billion a year earlier.[47]

The real bottleneck lies in operational transformation and not AI adoption. It is what the data across multiple research bodies, from global productivity statistics to enterprise-level surveys, now consistently confirms. The question for SMBs, and for the technology partners who serve them, is how to close the vast structural gap between surface-level adoption and the integrated, redesigned operating model that converts AI investment into compounding returns.

THE NUMBERS ARE GLOBAL AND THEY ARE NOT SUBTLE

The productivity case for agentic labor has been building across every major global research body, and what it reveals is a world caught between documented potential and uneven realization. The Organization for Economic Cooperation and Development (OECD) modeling projects that AI will raise labor productivity growth by between 0.5 and 3.5 percentage points per year across the G7 over the next decade; a range that, at its upper bound, would represent one of the most significant structural productivity shifts in modern economic history.[48] Yet the same OECD research documents that global labor productivity growth stagnated at approximately 0.4 percent on average across OECD member countries in 2024, and that the sizeable benefits of AI have yet to materially show up in aggregate productivity statistics at the global level.[49] The gap between those two numbers: the projected ceiling and the current floor, reveals the truth: this is a gap of adoption depth. The research is explicit on this point: realizing productivity gains from AI requires a combination of skilled labor, appropriate applications and complementary investments. Countries and businesses that have made those investments are beginning to see results.

The U.S. data offers a forward signal of what that shift looks like in practice. If 92 percent of SMBs are using digital tools and only 10 percent have fully integrated them, the 82 percent in between represent a market defined not by the absence of investment but by the absence of architecture. U.S. nonfarm labor productivity rose 2.8 percent year-over-year as of the most recent Bureau of Labor Statistics release; the highest single-year reading of the period that began when AI tools moved from experimentation into operational deployment at scale, and above even the 2.6 percent average annual rate this era has produced. The trend and the most recent data point are telling the same story from different angles: the productivity acceleration is real, it is continuing, and it is still building.[50] [51]

The mechanism connecting integration depth to productivity outcomes has been quantified at the industry level. Research published by the Federal Reserve Bank of St. Louis found that industries with higher AI-driven time savings are experiencing 2.7 percentage points faster productivity growth per point of time saved.[52] This is a documented relationship between agentic labor adoption and measured economic output, derived from real-time survey data tracking how workers across sectors are actually using AI in daily tasks. The industries leading the surge are not those with the most sophisticated enterprise infrastructures, but the ones in which AI-assisted workflows have moved from pilots into standard operating practice, where the tools are not supplementary but structural. That pattern is replicating across economies at different speeds, but the direction is consistent everywhere the research looks.

The mechanism connecting integration depth to productivity outcomes has been quantified at the industry level.

46 Sajadieh, Sha, et al. "The AI Index 2026 Annual Report." AI Index Steering Committee, Institute for Human-Centered AI, Stanford University, Apr. 2026.

47 Sajadieh, Sha, et al. "The AI Index 2026 Annual Report." AI Index Steering Committee, Institute for Human-Centered AI, Stanford University, Apr. 2026.

48 OECD. "OECD Compendium of Productivity Indicators 2025." OECD Publishing, Paris, 2025.

49 OECD. "OECD Compendium of Productivity Indicators 2025." OECD Publishing, Paris, 2025.

50 Bureau of Labor Statistics. "Productivity and Costs." U.S. Department of Labor, 2026.

51 Bureau of Labor Statistics. "Productivity Up 2.4 Percent in Second Quarter 2025." U.S. Department of Labor, 6 Sept. 2025.

52 Bick, Alexander, Adam Blandin, and David Deming. "The State of Generative AI Adoption in 2025." Federal Reserve Bank of St. Louis, 19 Nov. 2025.

The enterprise-level global surveys confirm the same pattern from a different angle. McKinsey's 2025 State of AI survey finds that 88 percent of organizations globally now use AI in at least one business function, up from 78 percent the prior year.[53] Deloitte's 2026 State of AI in the Enterprise report documents that 84 percent of organizations have increased their AI investments and 78 percent of leaders report growing confidence in the technology.[54] By conventional adoption measures, both surveys describe a world that has broadly committed to AI. But both surveys reach the same conclusion about what that commitment has and has not produced: most organizations are capturing productivity gains at the function or task level, while enterprise-wide financial impact remains concentrated among a small fraction of businesses that have gone further. McKinsey finds just 39 percent of organizations globally attribute any EBIT impact to AI, and only 7 percent have fully scaled AI across their organizations.[55] The adoption story is largely written. The integration story is only beginning.

What makes this data particularly significant for the SMB economy is the parallel adoption picture it reveals at the small business level. ASUS's 2025 Future of SMB Report finds that 53 percent of SMBs are already using AI, with 64 percent indicating readiness to adopt now and 41 percent already reporting tangible benefits.[51] By most conventional measures, this looks like a success story. More than half of small businesses have crossed the adoption threshold. But adoption, as the productivity data makes clear, is not the variable that explains the performance gap. What explains it is the depth of integration — and on that dimension, most small businesses have barely started.

THE LAST MILE PROBLEM

The distance between using AI and operationalizing it has a name in the research literature, and it is gaining traction as a precise description of where value is being lost. Harvard Business Review has described it as the "last mile" problem: the gap between the point at which an organization has deployed AI tools and the point at which those tools have been woven into the fundamental operating logic of the business.[56] What characterizes organizations stuck in this gap is not a lack of effort or investment. It is a structural pattern that McKinsey has documented across enterprise-level surveys: AI creates "islands of productivity," genuine, measurable improvements at the individual task or function level that, nonetheless, fail to translate into business outcomes because the processes, roles and governance structures surrounding those individuals have not been redesigned around the new capabilities.[57]

McKinsey's 2025 State of AI findings make this visible in the data. While 88 percent of organizations globally now use AI in at least one business function, only 7 percent have fully scaled AI across their organizations, and just 39 percent report any EBIT impact attributable to AI at the enterprise level.[58] Among the organizations that have broken through, the distinguishing characteristic "AI high performers," representing roughly 6 percent of respondents, is not the sophistication of the tools deployed; it is the extent to which they have fundamentally redesigned their workflows around AI capabilities. High performers are 2.8 times more likely than peers to report fundamental workflow redesign, and the intentional redesign of workflows carries one of the strongest statistical contributions to achieving measurable business impact of all the factors McKinsey tested.[59] For small businesses, the last mile problem is even more acute than it is for enterprises, precisely because the organizational scaffolding required to cross it; dedicated AI leadership, data infrastructure investment, governance frameworks, role redesign capacity; is harder to build at the SMB scale. The tools are accessible. The architecture is not.

WHAT THE DATA MEANS FOR THE BUSINESSES STILL DECIDING

The implication of the industry-level productivity correlation is precise and uncomfortable for businesses currently operating in the messy middle. The 2.7 percentage point productivity advantage per point of AI-driven time savings is not a one-time event. It is a compounding operational advantage that accumulates each quarter in which a business has structured its workflows around agentic labor while a competitor has not. Businesses that crossed the integration threshold eighteen months ago have had six quarters to build on that advantage. Their workflows are optimized. Their staff are proficient. Their data infrastructure is mature enough to support additional AI capability.

The businesses still deliberating have not merely deferred a benefit; they have allowed the gap to widen while the clock continued to run. When a sector's productivity growth rate accelerates, it is because a critical mass of businesses within that sector have made structural changes that enable more output per hour of labor. For any individual SMB that has not made those changes, the productivity data is an undeniable measurement of ground being ceded.

53 Singla, Alex, et al. "The State of AI in 2025: Agents, Innovation, and Transformation." McKinsey & Company, Nov. 2025.

54 Rowan, Jim, et al. "State of AI in the Enterprise." Deloitte AI Institute, Jan. 2026.

55 Singla, Alex, et al. "The State of AI in 2025: Agents, Innovation, and Transformation." McKinsey & Company, Nov. 2025.

56 Lakhani, Karim R., Jared Spataro, and Jen Stave. "The 'Last Mile' Problem Slowing AI Transformation." Harvard Business Review, 9 Mar. 2026.

57 Singla, Alex, et al. "The State of AI in 2025: Agents, Innovation, and Transformation." McKinsey & Company, Nov. 2025.

58 Singla, Alex, et al. "The State of AI in 2025: Agents, Innovation, and Transformation." McKinsey & Company, Nov. 2025.

59 Singla, Alex, et al. "The State of AI in 2025: Agents, Innovation, and Transformation." McKinsey & Company, Nov. 2025.

THE TECHNOLOGY PARTNER'S ROLE IN CLOSING THE GAP

There's a clear managed service opportunity for the channel: If 92 percent of SMBs are using AI tools and only 10 percent have fully integrated them, the 82 percent in between represent a market defined not by the absence of investment but by the absence of architecture. These businesses have already committed to digital transformation in principle. They have already incurred the initial costs of adoption. What they lack is the operational scaffolding: workflow redesign, data readiness, governance structures, staff enablement, that converts those investments into compounding returns.

If 92 percent of SMBs are using AI tools and only 10 percent have fully integrated them, the 82 percent in between represent a market defined not by the absence of investment but by the absence of architecture.

That scaffolding is precisely what a well-positioned technology partner delivers. And the data changes the terms of that conversation in a significant way. Technology providers who have historically framed their value proposition around cost savings or security risk management now have access to something more compelling: a documented macroeconomic shift in which their clients are either participating or falling behind. The data and insights found in this report are not vendor-generated ROI projections. They are quantifiable measurements of what is happening in sectors that have completed the journey from adoption to integration. For MSPs and MIPs framing their value proposition to SMB clients, the macro signal is the argument; one that positions the technology partner as the bridge between where most SMBs currently sit and where the data shows the high-performing businesses around them have already arrived.

Chapter 4 examines the economic architecture of that integration journey in detail: what the J-curve of digital labor ROI looks like across maturity levels, why many of the value lives in the second half of the adoption curve and why most SMBs stall at precisely the point where returns begin to accelerate.

88%

of organizations globally now use AI in at least one business function.

39%

report any EBIT impact attributable to AI at the enterprise level.

Part 02.

The Economics of Agentic Labor

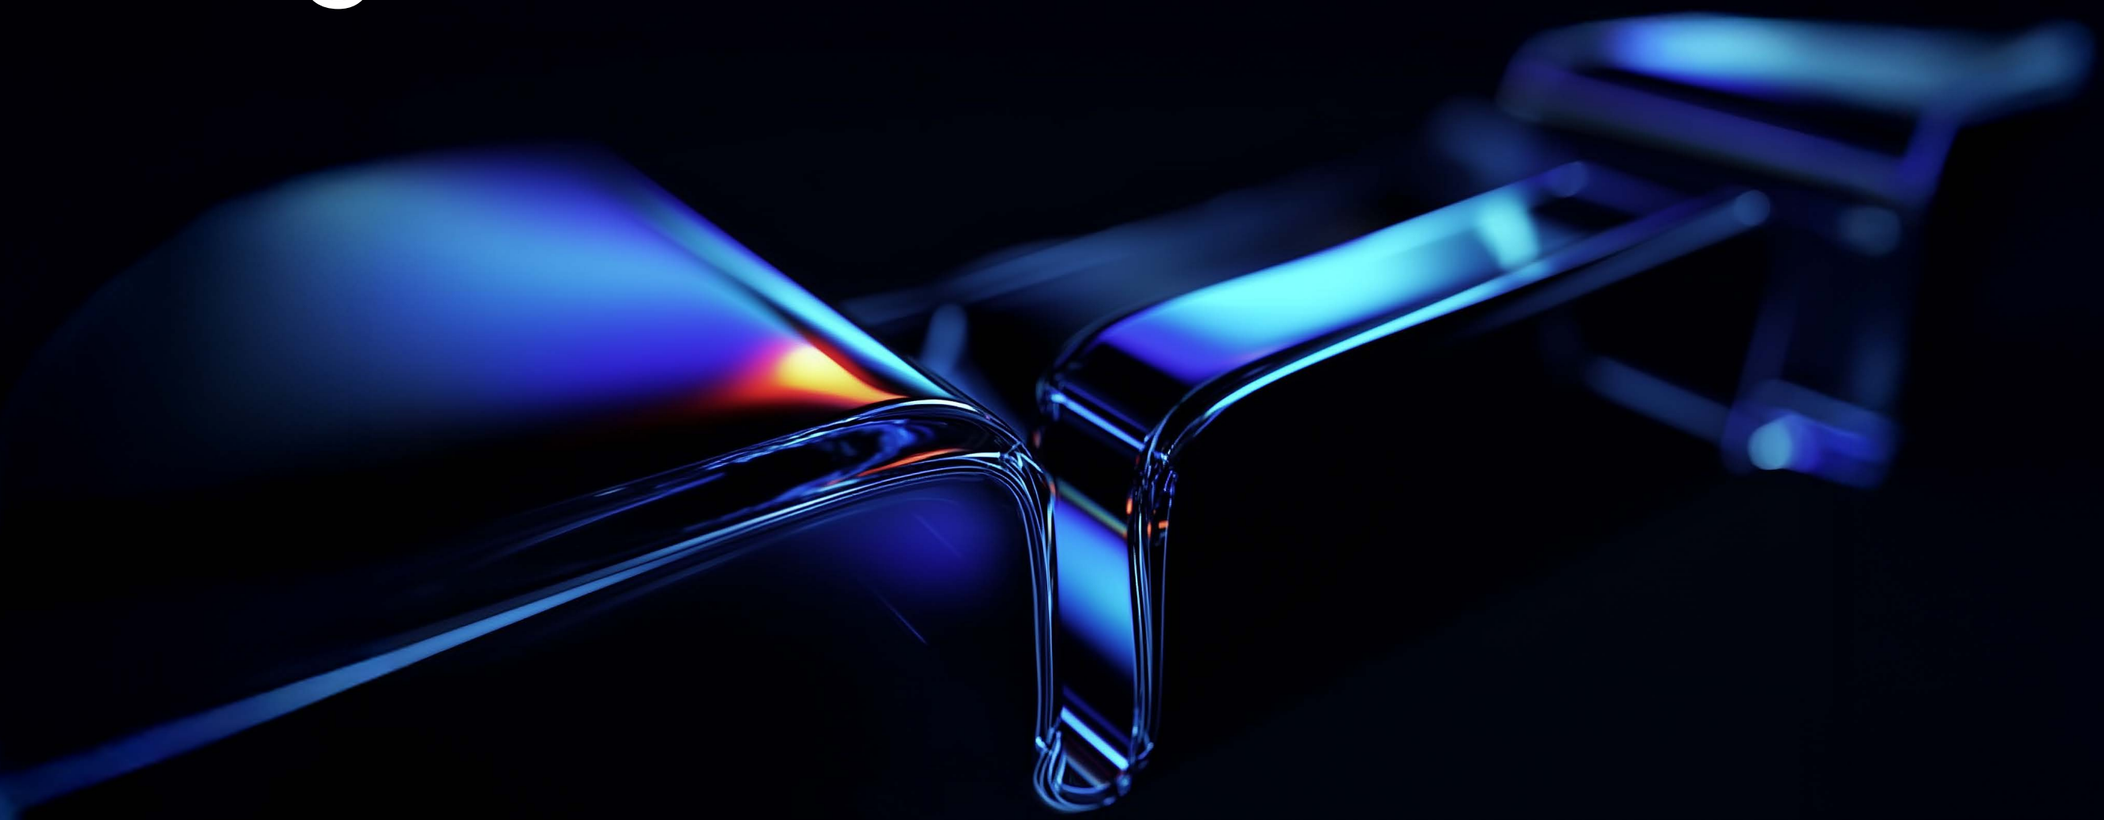

The adoption question is largely settled. Small businesses around the world are investing in AI tools at a rate that, eighteen months ago, would have seemed optimistic.

The adoption question is largely settled. Small businesses around the world are investing in AI tools at a rate that, eighteen months ago, would have seemed optimistic. The more consequential question, and the one this section of the report examines in detail, is what the economics of that investment look like across the maturity journey, and why the financial architecture of agentic labor adoption is so poorly understood by the businesses making it. The gap between the returns available to businesses that complete the integration journey and the returns being captured by those stuck at its early stages is a structural, compounding divide that reshapes competitive landscapes and determines which businesses survive the decade in which digital labor becomes the operating norm.

The chapters in Part Two examine four dimensions of that economic architecture: the nonlinear return curve that makes early adoption feel underwhelming and deep integration feel transformative; the time dividend that most SMBs are generating but few are reinvesting strategically; the leapfrog opportunity available to businesses with no legacy infrastructure to migrate; and the risk calculus that makes digital labor simultaneously the most productive investment and the most dangerous one an SMB can make without proper governance.

Chapter 04.

The J-Curve Most SMBs Never Reach

The empirical case for agentic labor's productivity impact is real, but it operates at the task level.

The economics of agentic labor adoption are widely misunderstood, and the misunderstanding is expensive. Most small businesses that have invested in AI tools and found the returns modest have not encountered a technology that overpromised. They have encountered a return curve that front-loads cost and back-loads value, and they have measured their results at the wrong point on it. The J-curve of technology adoption is not a new concept. Economists have applied it to electrification, the internet and enterprise software for decades. What makes it particularly consequential in the context of agentic labor is the steepness of the second phase: the uplift available to businesses that cross from intermediate to full integration is not incremental. It is, by Deloitte Access Economics modeling, roughly two and a half times larger than the gains captured in the first half of the journey.[60] Understanding why that curve is shaped the way it is, and where on it most SMBs currently sit, is the foundational economic insight that should drive every technology investment decision a small business makes over the next three years.

WHAT THE TASK-LEVEL EVIDENCE SHOWS

The empirical case for agentic labor's productivity impact is real, but it operates at the task level. Specific workers, doing specific things, in specific workflows designed or adapted to support AI assistance. And it does not automatically aggregate into business-level profitability improvement. That aggregation requires something the studies do not measure: the organizational redesign that converts scattered individual productivity gains into coordinated output at the company level. A business in which ten employees are each 15 percent more productive at their individual tasks; and the evidence for gains in that range is now robust, from customer service agents resolving issues faster, to developers completing coding tasks 55 percent faster, to less experienced workers seeing gains exceeding 30 percent as AI compensates for gaps in institutional knowledge; is not automatically a business that is 15 percent more profitable,[61] but rather, a business with unrealized potential, and whether that potential is realized depends entirely on what the organization does with the recaptured capacity. What makes this more urgent is that the gains are not static. Research tracking task-level professional productivity found that each successive generation of improved AI models reduces task completion time by approximately 8 percent; meaning the unrealized potential compounds as the technology matures.[62] Organizations that haven't built the infrastructure to capture today's gains will find themselves further behind with every model cycle.

60 O'Mahony, John, et al. "The AI Edge for Small Business: Increased SMB AI Adoption Can Add $44 Billion to Australia's Economy." Deloitte Access Economics, 25 Nov. 2025.

61 Peng, Sida, et al. "The Impact of AI on Developer Productivity: Evidence from GitHub Copilot." arXiv, Feb. 2023.

62 Siemon, Felix, et al. "AI Model Generations and Task Productivity." arXiv, Dec. 2025.

THE MATURITY LADDER AND THE NONLINEAR RETURN

Deloitte Access Economics has modeled the profitability implications of AI adoption across three maturity levels for SMBs, and the numbers illustrate the J-curve with unusual precision. Businesses that progress from basic to intermediate AI adoption see profitability uplifts of approximately 45 percent.[63] Those that progress from intermediate to fully enabled, fully integrated AI deployment see uplifts of approximately 111 percent.[64] The steepest part of the return curve is the second half of the journey, not the first.

The mechanism behind this nonlinearity is important to understand, because it explains both why early adopters often feel underwhelmed and why the businesses that push through to full integration find the investment transformative. In the first phase of the maturity journey, AI tools are being applied to individual tasks in existing workflows. The workflow itself has not changed. The roles around it have not changed. The data infrastructure supporting it has not changed. The gains are real but isolated; they show up in individual productivity metrics without flowing through to business-level margin improvement, because the organizational architecture around the tools has not been rebuilt to capture them.

In the second phase, the organizational architecture changes. Workflows are redesigned around AI capabilities rather than adapted to accommodate them. Roles are restructured to eliminate the human execution of tasks that AI can now handle reliably, redirecting human capacity toward judgment, relationship management and strategic work that AI cannot replicate. Data environments are structured and secured to give AI systems the clean, connected inputs they need to perform consistently.

The gains are real but isolated; they show up in individual productivity metrics without flowing through to business-level margin improvement, because the organizational architecture around the tools has not been rebuilt to capture them.

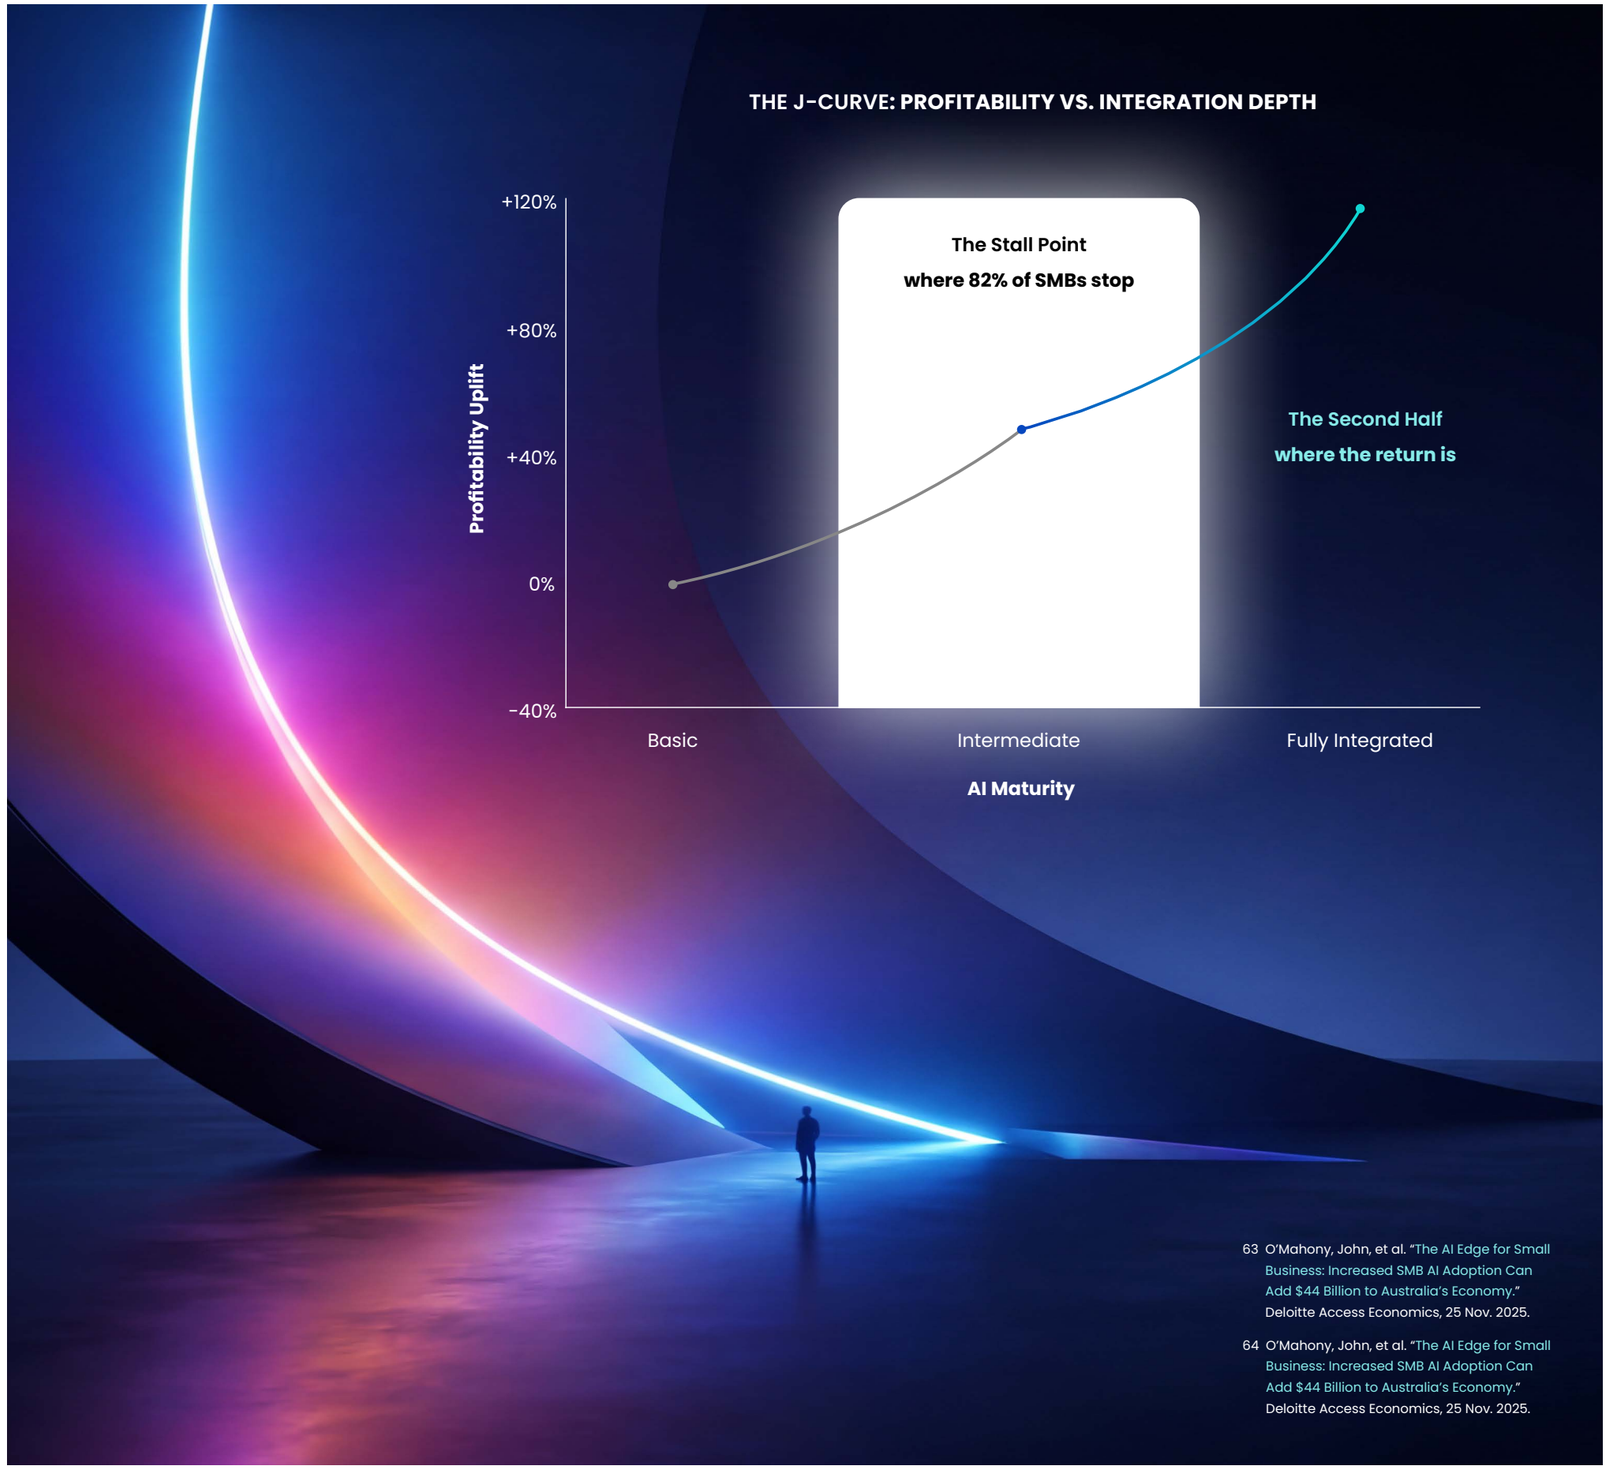

63 O'Mahony, John, et al. "The AI Edge for Small Business: Increased SMB AI Adoption Can Add $44 Billion to Australia's Economy." Deloitte Access Economics, 25 Nov. 2025.

64 O'Mahony, John, et al. "The AI Edge for Small Business: Increased SMB AI Adoption Can Add $44 Billion to Australia's Economy." Deloitte Access Economics, 25 Nov. 2025.

Governance frameworks are established that allow autonomous decision-making to proceed reliably within defined boundaries. When all these changes happen together, and they must happen together, because each one amplifies the others, the profitability impact is not additive. It is multiplicative. The 111 percent uplift at the fully enabled stage is not three times the 45 percent gain at the intermediate stage because three times as much AI is being used. It is because the organizational system around the AI has been built to convert AI output into business value at every point where the two intersect.

The correlation is consistent enough across multiple data sources to establish a clear behavioral pattern: SMBs that are scaling output without scaling payroll are disproportionately the ones investing in agentic labor, and they are doing so because the economics are compelling even before the full maturity curve is reached.

The organizational pattern this curve describes is now visible in aggregate productivity data. The most recent AI Index Report interprets the recent U.S. productivity acceleration as potentially reflecting the early stages of precisely this J-curve dynamic: a period in which organizations absorb the upfront costs of AI adoption before the larger returns compound into measurable macro-level gains.[65] The SMBs currently in the stall point between basic adoption and full integration are not experiencing a technology failure, but rather, the front half of a return curve whose back half is substantially larger.

THE STALL POINT AND WHY MOST SMBS NEVER LEAVE IT

The Canadian Federation of Independent Business documents the scale of the stall. 92 percent of SMBs are using digital tools. Only 10 percent have fully integrated them.[66] The 82-point gap between those figures is not primarily a technology gap. It is a stall point on the J-curve: a point at which the initial investment has been made, some task-level gains have been registered and the business has not yet committed to the organizational work required to reach the second phase of the return curve.

McKinsey's global survey data explains why the stall is so common. Among organizations at every scale, the transition from piloting to scaling AI is where most businesses get stuck.

Only one-third of respondents globally say their organizations are actively scaling AI, and just 7 percent have fully deployed it across their operations.[67] The barrier is not confidence in the technology; 88 percent of organizations are using AI in at least one function, and awareness of its potential has never been higher.[68]

The barrier is the organizational investment required to move from isolated deployment to integrated architecture. That investment is harder, slower and less visible than the initial tool purchase. It requires cross-functional coordination, senior leadership commitment, data infrastructure work and a willingness to redesign roles and processes that have been stable for years. For a small business owner managing day-to-day operations without a dedicated technology team, it is precisely the kind of work that gets deferred in favor of more immediate priorities; and each deferral pushes the payoff further out while the compounding gap with competitors who have crossed the threshold continues to widen.

HBR's analysis of the "last mile" problem captures the organizational character of this stall with precision. The businesses stuck between adoption and integration are organizations that have moved confidently through the first phase of the journey and encountered the genuine difficulty of the second: the point at which AI meets operating model design, and at which the barriers are no longer technical but structural.[69] Process fragmentation, governance bottlenecks, role redesign requirements and cultural resistance to workflow change are the friction that keeps the 82 percent from becoming the 10 percent.

There is a name for what the stall point produces at the market level: the speed gap. The relevant divide in the SMB economy is no longer between businesses that have access to AI technology and those that do not. The tools are available, affordable and increasingly embedded in the software most small businesses already pay for.

The divide is between the businesses iterating rapidly; deploying, learning, adjusting and redeploying; and the businesses that adopted once and stopped. The speed of iteration is now the primary competitive variable. A business that deploys one agent, learns from its performance and deploys a second based on that learning is one full learning cycle ahead, and that cycle advantage compounds with every subsequent deployment in a way that the laggard cannot close by deploying faster later. The J-curve rewards the businesses that complete it first and keep going. Technology partners who understand this are not selling software to businesses that have stalled. They are offering the organizational architecture that converts stalled investment into compounding returns.

The speed of iteration is now the primary competitive variable.

WHAT THE HIGH PERFORMERS DO DIFFERENTLY

McKinsey's identification of AI high performers; the roughly 6 percent of organizations globally that attribute more than 5 percent of EBIT to AI and report significant enterprise-wide value; provides the clearest available picture of what crossing the J-curve actually requires.[70] These organizations are not distinguished primarily by the tools they use or the scale of their AI budgets. They are distinguished by a set of organizational practices that consistently predict value realization across every sector and business size studied.

65 Sajadieh, Sha, et al. "The AI Index 2026 Annual Report." AI Index Steering Committee, Institute for Human-Centered AI, Stanford University, Apr. 2026.

66 Alegbeh, Alchad, and Marvin Cruz. "Digital Transformation: How Small Businesses in Canada Are Leveraging AI and Technology for Growth and Productivity." Canadian Federation of Independent Business, 29 Sept. 2025.

67 Singla, Alex, et al. "The State of AI in 2025: Agents, Innovation, and Transformation." McKinsey & Company, Nov. 2025.

68 Singla, Alex, et al. "The State of AI in 2025: Agents, Innovation, and Transformation." McKinsey & Company, Nov. 2025.

69 Lakhani, Karim R., Jared Spataro, and Jen Stave. "The 'Last Mile' Problem Slowing AI Transformation." Harvard Business Review, 9 Mar. 2026.

70 Singla, Alex, et al. "The State of AI in 2025: Agents, Innovation, and Transformation." McKinsey & Company, Nov. 2025.

HIGH PERFORMERS DISTINGUISH THEMSELVES FROM THEIR PEERS IN FOUR DISTINCT WAYS:

- **They redesign workflows:** High performers are 2.8 times more likely to have fundamentally redesigned their workflows around AI, and the relative weights analysis identifies this workflow redesign as carrying one of the strongest statistical contributions to achieving measurable business impact of all the factors tested.[71]

- **Their leaders lead by example:** Those at the top are 3 times more likely to have senior leaders who demonstrate active ownership of and commitment to AI initiatives, including role modeling the use of AI across the organization.[72]

- **They build governance that enables autonomy:** High performers are significantly more likely to have defined processes for determining when model outputs require human validation; a governance discipline that enables autonomous AI decision-making to proceed reliably rather than requiring constant manual oversight.[73]

- **They pursue growth alongside efficiency:** This cohort are more likely to have set growth and innovation as explicit objectives of their AI programs, alongside efficiency; a framing that McKinsey finds consistently predicts a wider range of qualitative enterprise benefits, from competitive differentiation to customer satisfaction improvement.[74]

The IDC study of frontier firms adds a further dimension to this picture. Among the organizations generating the highest returns from AI globally, the differentiating characteristic is the decision to pursue full-stack integration; combining data infrastructure, technology deployment, human expertise and governance architecture into a unified operating system rather than treating each element as a separate initiative.[75] IDC finds that frontier firms generate approximately 3 times the ROI from their AI investments that laggard organizations generate, with the gap driven almost entirely by integration depth rather than by investment volume.[76] **Spending more on AI tools without closing the integration gap does not move a business up the curve. Building the organizational architecture around those tools does.**

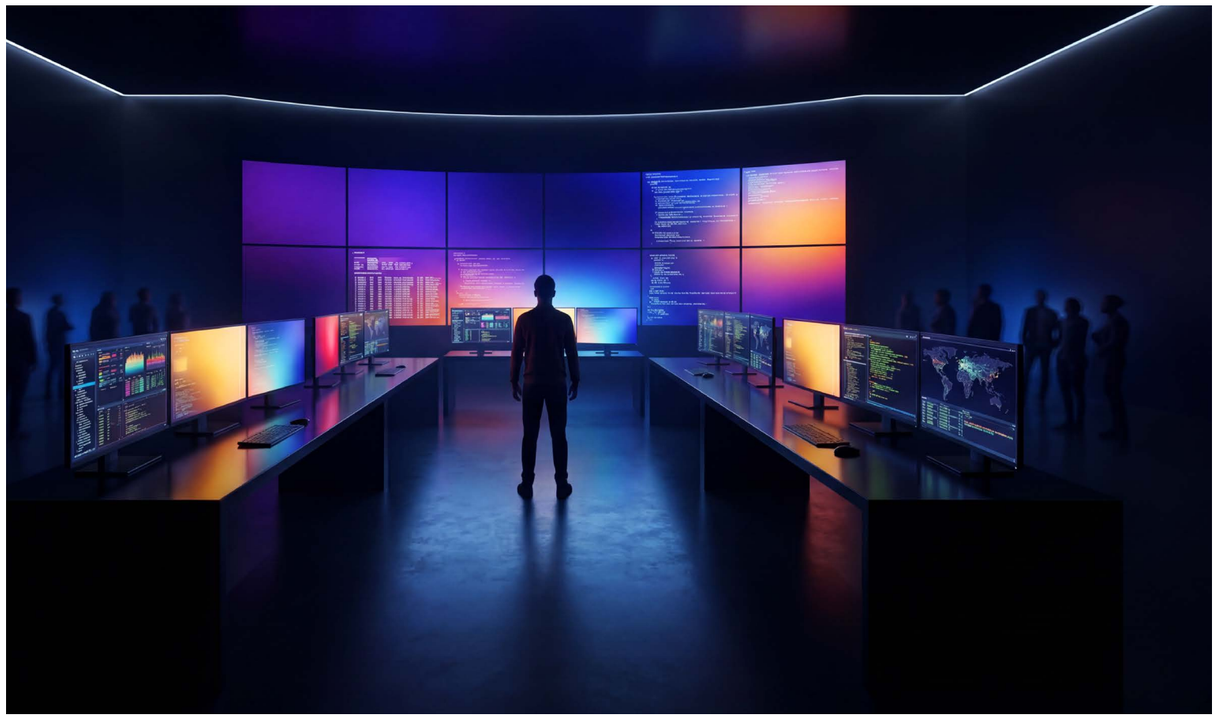

THE TECHNOLOGY PARTNER'S ROLE IN THE SECOND HALF OF THE CURVE

The economic architecture of the J-curve has a direct implication for how technology partners should position their value proposition, and it is an implication that most of the channel has not yet fully internalized. The first phase of the digital labor maturity journey, in which individual AI tools are deployed for specific tasks in existing workflows, is a phase characterized by relatively low margin, high competitive pressure and modest switching costs.

Any number of providers can help an SMB adopt a generative AI subscription, configure an automated workflow or deploy a chatbot for customer intake. The value delivered is real but not defensible.

The second phase is different in every dimension. A technology partner helps an SMB redesign its core operating model around AI capabilities: rebuilding workflows, rationalizing data environments, restructuring roles, establishing governance frameworks and sequencing the integration of AI agents across interconnected business processes. The work is complex, the expertise required is scarce and the switching costs compound with every milestone reached. The client who has rebuilt their operations around an agentic workforce with a trusted partner has built a competitive advantage, and that advantage is inseparable from the partner who helped architect it. The 111 percent profitability uplift at the fully enabled stage is not only the return available to the SMB, but the value that the technology partner makes possible and, in doing so, makes themselves indispensable to.

The businesses and technology partners that understand this, understand the fundamental economics of the decade ahead. The J-curve is a value architecture to be navigated and the partners who guide clients through its second half will define the category.

The client who has rebuilt their operations around an agentic workforce with a trusted partner has built a competitive advantage, and that advantage is inseparable from the partner who helped architect it.

Chapter 5 looks at the most immediate and most squandered output of the first phase of the maturity journey: the time that digital labor is already recapturing for SMB workers, and the striking finding that fewer than one in five businesses has a deliberate strategy for where that time goes.

71 Singla, Alex, et al. "The State of AI in 2025: Agents, Innovation, and Transformation." McKinsey & Company, Nov. 2025.

72 Singla, Alex, et al. "The State of AI in 2025: Agents, Innovation, and Transformation." McKinsey & Company, Nov. 2025.

73 Singla, Alex, et al. "The State of AI in 2025: Agents, Innovation, and Transformation." McKinsey & Company, Nov. 2025.

74 Singla, Alex, et al. "The State of AI in 2025: Agents, Innovation, and Transformation." McKinsey & Company, Nov. 2025.

75 Evans, Katie, et al. "The SMB 2026 Digital Landscape: How AI Is Redefining Growth." IDC, 10 Feb. 2026.

76 Evans, Katie, et al. "The SMB 2026 Digital Landscape: How AI Is Redefining Growth." IDC, 10 Feb. 2026.

Chapter 05.

The Three-Hour Dividend Nobody Is Reinvesting

Every conversation about agentic labor's impact on small business productivity eventually arrives at a number that commands attention. IDC's research on the digital labor economy finds that line-of-business (LOB) workers using AI tools save approximately 3.1 hours per day, roughly 39 percent of the standard workday.[77] IT workers save even more: approximately 3.6 hours daily, or 45 percent of their working hours.[78] These are the documented, recurring savings being generated right now, across hundreds of thousands of businesses that have adopted AI tools with even modest consistency.

The number that should command equal attention, but rarely does, is the one that follows: fewer than one in five SMBs have a deliberate plan for where that time goes.

The three-hour dividend is one of the most significant and most squandered assets in the small business economy. It is being generated daily, at scale, by businesses that have already made the tool investments required to produce it. And in the overwhelming majority of cases, it is evaporating; absorbed by unstructured work, redistributed across existing tasks without any strategic direction or simply not accounted for in any organizational planning conversation. The businesses that treat recaptured time as a strategic resource rather than a pleasant side effect of their technology spend are the ones seeing compounding returns. The businesses treating it as freed capacity without a destination are performing the small business equivalent of generating revenue and leaving it in a non-interest-bearing account.

Part of what suppresses deliberate reinvestment is not strategic oversight but human caution. When Pax8's research team interviewed SMB owners and operational leaders about their relationship with AI, the most consistently named concern was not data quality or cost or competitive disadvantage. It was the fear that aggressive AI deployment would amount to training the technology to replace the people who currently operate the business. The concern was most acute among small, traditionally labor-intensive operators (blue-collar trades, family-owned services, owner-led creative practices) where the recaptured hour is more easily framed as a threat to a livelihood than as capacity to redirect.[79] That framing has a direct bearing on the reinvestment problem. A business owner who is uncertain whether the recovered time is meant to expand the business or eliminate roles within it tends to leave the question unanswered, and unanswered time becomes absorbed time. The reinvestment plan is missing in part because the conversation it requires (about what the business is for, who it employs and how AI changes both) has not yet happened.

THE THREE-HOUR DIVIDEND. TIME RECAPTURED VS. TIME REINVESTED

WHAT THE SAVINGS LOOK LIKE IN PRACTICE

The IDC figures capture time savings at the category level, but the granular data across multiple research bodies reveals what those savings look like inside actual small business operations. LinkedIn's December 2025 Work Change Special Report found that the average SMB using generative AI saved 1.08 hours per working day; more than an hour of reclaimed capacity, every day, in businesses where every hour carries disproportionate weight relative to available headcount.[80] The CFIB's research on Canadian SMBs documents productivity gains of 29 percent among businesses that have moved AI tools into regular operational use, with a $1.60 return documented for every dollar invested in digital tools across the sample.[81]

These returns are being felt most acutely in the operational functions that consume SMB owner and employee time without generating proportional value: reconciling accounts, triaging incoming requests, scheduling and rescheduling, generating routine reports, following up on outstanding payments.[82] These are functions that require consistency and benefit from scale, precisely the conditions under which AI agents outperform human labor on a cost-per-unit basis. When an AI agent handles appointment confirmations, invoice follow-up and intake processing, the hours those tasks previously consumed do not disappear from the business. They are returned to the humans who were spending them, and what happens to those hours next is the central economic question this chapter examines.

The CFIB's research on Canadian SMBs provides additional texture on where the perceived value is concentrated among current users. The highest-rated AI benefits are improved data analysis and decision-making, better planning and forecasting capability, and the ability to summarize information and eliminate repetitive tasks.[83] These correspond directly to the daily operational frictions that consume disproportionate time in small businesses across every industry: the hours spent assembling information that should be aggregated automatically, the time spent on follow-up sequences that a well-configured agent could handle end-to-end, the cognitive overhead of switching between fragmented administrative tasks that software could queue and execute without human intervention.

These returns are being felt most acutely in the operational functions that consume SMB owner and employee time without generating proportional value.

77 Railton, Matt. "IDC Defines Digital Labor for the Agentic Enterprise." No Jitter, 15 Oct. 2025.

78 Railton, Matt. "IDC Defines Digital Labor for the Agentic Enterprise." No Jitter, 15 Oct. 2025.

79 Galvan, Moriah. "SMB Discovery." Pax8 UX Research, May 2026. Proprietary research. Data on file.

80 LinkedIn Economic Graph. "How Small Businesses Can Win in 2026: Work Change Special Report." LinkedIn, Dec. 2025.

81 Alegbeh, Alchad, and Marvin Cruz. "Digital Transformation: How Small Businesses in Canada Are Leveraging AI and Technology for Growth and Productivity." CFIB, 29 Sept. 2025.

82 Railton, Matt. "IDC Defines Digital Labor for the Agentic Enterprise." No Jitter, 15 Oct. 2025.

83 Alegbeh, Alchad, and Marvin Cruz. "Digital Transformation: How Small Businesses in Canada Are Leveraging AI and Technology for Growth and Productivity." CFIB, 29 Sept. 2025.

THE STRANDED DIVIDEND
The problem is not that SMBs are failing to generate time savings. Many are generating them consistently. The problem is that without a deliberate plan for redirecting recaptured hours, those savings are functionally invisible to the business's bottom line. Time that is not deliberately reallocated tends to be absorbed by the ambient demands of day-to-day operations; the inbox that refills, the client calls that expand to fill available space, the low-priority tasks that surface to occupy gaps in the schedule. The productivity gain is real at the individual level; it simply fails to aggregate into business-level returns because the organization has not made a structural decision about where the recaptured capacity goes.

This is the distinction that separates the businesses seeing compounding returns from those seeing modest efficiency improvements. The J-curve examined in Chapter 4 describes the nonlinear return structure of the full integration journey. The three-hour dividend is where that curve is either accelerated or stalled at the earliest stage. A business that recovers 3.1 hours per LOB worker per day and immediately redeploys those hours toward customer acquisition, product development or service quality improvement is on the accelerating portion of the curve. A business that recovers the same hours and allows them to dissolve back into unstructured work has generated a productivity statistic without generating a business result.

McKinsey's high-performer research explains the dynamics with regards to high performers at the enterprise level. The organizations generating measurable EBIT impact from AI are the ones that have embedded AI into redesigned workflows and made explicit decisions about how recaptured human capacity is directed.[84] High performers are 2.8 times more likely to have fundamentally redesigned their workflows, a practice that includes not only restructuring how tasks are performed but determining what humans do with the time that workflow automation returns to them.[85] The time savings and the reinvestment decision are, fundamentally, two stages of the same strategic process, and the gap between them is where most of the value is currently being lost.

THE COMPOUNDING RETURN OF DELIBERATE REINVESTMENT
The CFIB's modeling of the GDP impact of digital tool adoption among Canadian SMBs offers one of the most direct available quantifications of what deliberate time reinvestment looks like at scale. The research models the economic impact of SMBs reinvesting even a portion of their AI-generated time savings into higher-value activities, and the aggregate effect is significant: an indication that the mechanism is not simply about individual business efficiency but about the cumulative economic contribution of small businesses operating at a higher output-per-hour baseline.[86]

At the individual business level, the compounding dynamic works through a specific sequence. In the first phase, AI tools recover hours from administrative and operational tasks. In the second phase, those hours are redirected; toward customer relationships that generate retention, toward service quality improvements that generate referrals, toward product or service development that generates new revenue, or toward the organizational work of building the data foundations and governance frameworks required for the deeper integration that Chapter 4 identified as the source of the 111 percent profitability uplift. Each of these reinvestment decisions creates a return that compounds over time. The customer retained generates recurring revenue. The referral generates an acquisition at zero cost. The service improvement generates a premium pricing opportunity. The data foundation enables AI capabilities that generate further time savings. The business that manages this reinvestment cycle deliberately is on a compounding advantage trajectory. The one that allows the hours to dissolve is running in place.

THE REINVESTMENT GAP AS A PARTNER OPPORTUNITY
The finding that fewer than one in five SMBs have a deliberate reinvestment plan for their AI-generated time savings is a precise description of where a technology partner can deliver value that no software tool can automate. Deploying an AI agent that saves a client fifteen hours per week is Phase 1 work. Facilitating the structured conversation that determines where those fifteen hours go: which strategic priorities they fund, which growth activities they enable, how they are tracked and reported as a return on the AI investment; is Phase 2 work, and it is the kind of work that converts a technology vendor relationship into an outcome partnership.

The MSP or MIP who builds a reinvestment planning framework into every agentic labor deployment is doing something categorically different from the provider that measures success by deployment completion. They are providing the organizational architecture that converts tool-level savings into business-level returns, and in doing so, they are creating the measurable ROI story that justifies renewal, expansion and referral. The three-hour dividend is already being generated by clients across the channel. The providers who help clients capture it strategically will be the ones those clients associate with the results.

Chapter 6 examines a parallel economic advantage available specifically to the smallest businesses in the SMB economy; one that inverts the traditional assumption that larger businesses are further along the AI adoption curve, and that turns the absence of legacy infrastructure from a perceived disadvantage into a structural competitive edge.

The MSP or MIP who builds a reinvestment planning framework into every agentic labor deployment is doing something categorically different.

84 Singla, Alex, et al. "The State of AI in 2025: Agents, Innovation, and Transformation." McKinsey & Company, Nov. 2025.

85 Singla, Alex, et al. "The State of AI in 2025: Agents, Innovation, and Transformation." McKinsey & Company, Nov. 2025.

86 Alegbeh, Alchad, and Marvin Cruz. "Digital Transformation: How Small Businesses in Canada Are Leveraging AI and Technology for Growth and Productivity." CFIB, 29 Sept. 2025.

Chapter 06.

The AI-Built SMB: The Leapfrog Advantage of Starting Clean

Every prior technology adoption cycle in the history of small business followed the same curve. Larger businesses adopted first, absorbing the early costs and complexity while building internal capability. Smaller businesses watched, waited and followed, sometimes years later, sometimes a full business generation behind. Broadband internet reached enterprise before it reached the corner store. Cloud computing was a large-organization capability for the better part of a decade before it became a small-business default. Mobile payments, e-commerce infrastructure, CRM software: the pattern repeated itself with remarkable consistency across every major platform shift of the past thirty years.

That curve is inverting. And for the smallest businesses in the economy; the micro-SMBs, the startups, the businesses being built right now in industries that have always been labor-intensive and infrastructure-light; the inversion is a structural competitive edge that the mid-market organizations they will compete against cannot easily replicate, because that which those mid-market organizations spent years accumulating is now slowing them down.

These businesses entering the market today with no legacy systems to migrate, no entrenched processes to reengineer and no organizational habits to overcome are "AI-built SMBs." These are businesses that do not adopt AI onto existing operations. They build their operations around AI from the first day, and in doing so, they achieve a level of operational efficiency and scalability that would cost an established competitor months of disruptive, expensive transformation work to replicate.

AI-built SMBs build their operations around AI from the first day.

THE GAP THAT IS CLOSING AND WHAT IT REVEALS

The speed of the adoption gap's convergence is the clearest available signal that the traditional technology adoption curve has broken down. The U.S. Census Bureau Business Trends and Outlook Survey data documents the compression in real time: in early 2024, large businesses used AI at 1.8 times the rate of small businesses.[87] By August 2025, that ratio had narrowed to approximately 1.2 times.[88] Small businesses may now be only a year behind their larger counterparts; a gap that in prior technology cycles measured in years or decades.[89]

The IDC FutureScape data illustrates the demand-side driver of that compression. The share of SMBs listing AI as a top technology investment priority jumped from 16 percent in 2023 to 35 percent in 2024; a doubling in a single year, reflecting the transition of AI from a subject of curiosity to a strategic imperative across the small business economy.[90] When a technology priority doubles in twelve months, the adoption curve does not simply accelerate, but restructures, as the population moving toward adoption shifts from early adopters to mainstream buyers and the tools available to that mainstream buyer improve to meet them.

A recent Microsoft report captured the destination of that convergence with a data point that should reframe how the channel thinks about market segmentation. Overall, AI and generative AI adoption among SMBs stands at 71 percent. Among digital-native firms, businesses built on digital infrastructure from inception, adoption reaches 90 percent.[91] That 19-point gap between the overall SMB population and its digital-native subset is the born-agentic advantage made visible.

Businesses that started clean did not need a migration roadmap. They were already there.

Salesforce's global survey of SMB leaders adds the performance dimension that explains why that gap is not closing on its own. 91 percent of SMBs using AI report that it directly boosts their revenue.[92] But the more consequential finding is the divergence in forward investment intent between growing and declining SMBs: 78 percent of growing businesses plan to increase AI investment in the coming year, compared to 55 percent of declining businesses.[93] The businesses succeeding with AI are doubling down. The businesses struggling are pulling back. Each quarter that cycle repeats, the performance gap between the two cohorts widens, and the legacy SMB that defers its transformation journey is not simply staying even with the born-agentic competitor entering its market. It is falling further behind a business that started with a structural advantage and is compounding it.

THE LEGACY DEBT THAT BECOMES A CEILING

The central insight of this chapter is counterintuitive enough to warrant careful statement. For most of the history of technology adoption, size and incumbency were advantages. Larger businesses had more resources to invest, more organizational capacity to absorb disruption and more existing infrastructure to build on. In the current transition, a specific component of that advantage, existing infrastructure, has become a liability rather than an asset, and the businesses that lack it are discovering that their disadvantage has quietly become an edge.

Legacy infrastructure is expensive to maintain, slow to migrate and fundamentally resistant to the kind of clean-slate architectural redesign that AI-first operating models require. A mid-market business that built its operations around an on-premises ERP system, a suite of point-solution SaaS tools accumulated over a decade and a workforce trained on manual workflows cannot simply overlay AI agents onto that architecture and expect the compounding returns of AI leaders. Bain's recent Technology Report frames the stakes for established organizations with unusual directness: AI leaders have moved from pilots to profits, delivering 10 to 25 percent EBITDA gains by scaling AI across core workflows, while organizations that are still in the piloting phase are, in Bain's characterization, dangerously behind.[94] The integration complexity compounds with every legacy system in the stack. The change management required to redesign workflows around AI capabilities encounters resistance at every point where existing habits and role identities have been built around the processes that AI is now replacing.

87 Alegbeh, Alchad, and Marvin Cruz. "Digital Transformation: How Small Businesses in Canada Are Leveraging AI and Technology for Growth and Productivity." Canadian Federation of Independent Business, 29 Sept. 2025.

88 U.S. Small Business Administration, Office of Advocacy. "AI in Business: Small Firms Closing In." SBA, Sept. 2025.

89 U.S. Small Business Administration, Office of Advocacy. "AI in Business: Small Firms Closing In." SBA, Sept. 2025.

90 Evans, Katie, et al. "The SMB 2026 Digital Landscape: How AI Is Redefining Growth." IDC, 10 Feb. 2026.

91 Microsoft. "Microsoft's Fifth Annual Small and Medium-Sized Business Report." Microsoft, 25 June 2025.

92 Salesforce. "New Research Reveals SMBs with AI Adoption See Stronger Revenue Growth." Salesforce, 2025.

93 Salesforce. "New Research Reveals SMBs with AI Adoption See Stronger Revenue Growth." Salesforce, 2025.

94 Bain & Company. "State of the Art of Agentic AI Transformation." Bain Technology Report 2025, 2025.

The born-agentic business does not have to make that transition. It starts where the transition leads.

A born-agentic business has none of those constraints. It can select a cloud-native accounting platform designed from the ground up to support AI-driven reconciliation. It can implement a CRM that connects natively with AI agents for follow-up and customer communication. It can build its intake, scheduling and client management workflows around automation from the first week of operation, before any manual habit has formed or any employee's professional identity has been built around a process that an agent will eventually handle.

Gartner's prediction that 20 percent of organizations will use AI to flatten their organizational structures in 2026, eliminating more than half of current middle management positions as AI absorbs the coordination and oversight functions those layers previously performed, illuminates another dimension of the born-agentic advantage.[95] An established business flattening its structure mid-operation faces the disruption of eliminating roles that people have built careers around, restructuring reporting relationships that have been stable for years and managing the cultural friction of a reorganization. A born-agentic business simply does not build the hierarchical layers in the first place. It begins with the flat, AI-coordinated structure that its legacy competitors are trying to migrate toward, and it builds with it rather than against it.

The right historical analogy is factory electrification. When electricity first arrived in the late nineteenth century, many manufacturers simply replaced the steam engine with an electric motor while preserving the old line-shaft architecture; a single belt running through the entire facility, driving every machine from one central power source. The source of energy changed, but the operating model did not. The real productivity gains arrived a generation later, when factories were redesigned from the ground up around distributed electric drive: machines placed where the work required them rather than where the belt could reach them, workflows organized around output logic rather than power transmission geometry and the entire physical architecture of production rebuilt around what the new energy source actually made possible rather than what the old one had required.

AI adoption in the small business economy is following the same pattern. The first wave inserts AI into legacy workflows, a new capability grafted onto an organizational architecture built for a different era. The real gains will come in the second wave, when businesses redesign their operations around what agents can do: persistent execution, continuous monitoring, autonomous handling of rule-governed work and human attention reserved for judgment, relationships and decisions that genuinely require it. The born-agentic business does not have to make that transition. It starts where the transition leads. Its workflows were never built around the line-shaft model, so it never has to tear one out.

WHAT BORN-AGENTIC LOOKS LIKE IN PRACTICE

The abstract case for the born-agentic advantage becomes concrete when examined through the lens of businesses already living it. Forbes has documented examples across industries that illustrate the pattern with striking clarity: Dan Peguine, helping transform his parents' tea business, deployed AI agents to manage ordering, inventory and vendor communications — estimating annual savings of $50,000 and reclaiming two to three days each month.[96] Nathan Baschez, founder of Lex, an AI-native writing and editing platform, reports that working with AI agents helps him move 15 to 20 times faster; a pace of execution that compounds into competitive advantage the longer the gap between adopters and non-adopters widens.[97]

Pax8's own discovery research surfaced a comparable example at the operational level. Among 18 interviewed SMBs, one IT administrator at a small business had automated their internal IT operations with AI tools to the point of concluding that the business no longer had a meaningful use case for a managed service provider. The decision was not framed as a practical observation that the operational work an MSP would deliver was already being executed by the tools the business had configured.[98] That outcome (one micro-organization arriving, through individual initiative, at the operating posture this chapter has been describing) reinforces the chapter's structural point. The born-agentic advantage is available to any business willing to build its operations around AI from the outset, and the businesses that recognize this earliest will redefine what their size category is capable of.

The born-agentic advantage is available to any business willing to build its operations around AI from the outset, and the businesses that recognize this earliest will redefine what their size category is capable of.

The workforce model that born-agentic businesses are developing is equally distinctive. Wired's reporting on the emerging category of AI-native operators documents what this looks like at the organizational level: Simon Last, a cofounder of Notion, now manages multiple AI agents rather than writing code himself: a fundamental shift in what "doing the work" means for a founding team member.[99] Companies like Notion are hiring fewer people but prioritizing individuals who excel at directing and managing AI systems, creating a workforce composition that is structurally leaner and operationally more capable than traditional headcount-equivalent teams. The born-agentic SMB is not simply a smaller version of a traditional business with AI tools added. It is a different organizational model, one in which human expertise is concentrated at the orchestration and judgment layer while agents handle execution at scale.

This is the operating model that IDC's research on the digital labor economy describes as the direction of travel for the entire SMB economy but born-agentic businesses are not traveling toward it. They arrived there at inception.[100]

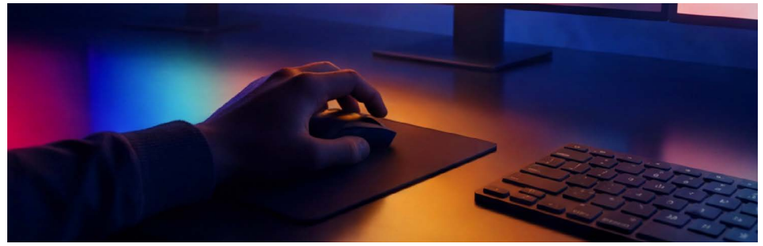

THE SMB GROUP MATURITY PICTURE AND THE DIVERGING TRAJECTORIES

Research on digital transformation stages among small and medium businesses maps the population distribution across maturity levels and documents the performance differential between them.[101] The finding most relevant to this chapter is the distribution itself: a significant majority of SMBs remain in the early and intermediate stages of the transformation journey, with the share at the highest maturity level, full integration across all operational functions, remaining small.[102]

For the legacy SMBs in those early stages, the path to full integration runs through the difficult, expensive, time-consuming work of migrating existing processes, retraining existing staff and rearchitecting existing data environments. The Omdia research cited by Canalys finds that 61 percent of channel partners struggle to move AI projects beyond proof of concept with their existing clients; a statistic that reflects the genuine difficulty of the migration journey rather than any failure of ambition.[103] For born-agentic businesses starting clean, there is no proof-of-concept phase to navigate because there are no legacy workflows to migrate. The question is not how to move from current state to desired state. It is simply what the desired state looks like and then building it.

McKinsey's global data confirms that the performance differential between maturity levels applies equally to both cohorts, but the time required to traverse the distance is systematically shorter for businesses with no legacy debt. Among organizations below $100 million in revenue, 29 percent have reached the AI scaling phase, compared to 49 percent of organizations above $5 billion.[104] The born-agentic business begins much closer to that scaling threshold than a legacy SMB at the same revenue level, because the starting point of its architecture is further along the maturity curve.

95 Gartner. "Gartner Identifies the Top 10 Strategic Technology Trends for 2025." Gartner, 21 Oct. 2024.

96 Griffin, Anne. "AI Agents For Business Are Creating a Gap Few Companies See Coming." Forbes, 25 Feb. 2026.

97 Griffin, Anne. "AI Agents For Business Are Creating a Gap Few Companies See Coming." Forbes, 25 Feb. 2026.

98 Galvan, Moriah. "SMB Discovery." Pax8 UX Research, May 2026. Proprietary research. Data on file.

99 Zeff, Maxwell. "Are You 'Agentic' Enough for the AI Era?" Wired, 26 Feb. 2026.

100 Evans, Katie, et al. "The SMB 2026 Digital Landscape: How AI Is Redefining Growth." IDC, 10 Feb. 2026.

101 McCabe, Laurie, and Sanjeev Aggarwal. "The Impact of AI on SMBs: 2025 Trends, Challenges, and Opportunities." SMB Group, June 2025. Sponsored by Workday.

102 McCabe, Laurie, and Sanjeev Aggarwal. "The Impact of AI on SMBs: 2025 Trends, Challenges, and Opportunities." SMB Group, June 2025. Sponsored by Workday.

103 Evans, Katie, et al. "The SMB 2026 Digital Landscape: How AI Is Redefining Growth." IDC, 10 Feb. 2026.

104 Singla, Alex, et al. "The State of AI in 2025: Agents, Innovation, and Transformation." McKinsey & Company, Nov. 2025.

THE COMPOUNDING ADVANTAGE THE RESEARCH CONFIRMS

The performance gap between AI adopters and non-adopters documented across the research base is not a static divide. It is a compounding curve, and Salesforce's finding that growing businesses are investing in AI at a 42-percent higher rate than declining ones is the mechanism through which that compounding accelerates.[105] The growing business that is already generating AI-driven returns reinvests those returns into deeper AI capability. The declining business that is deferring AI investment to focus on short-term cost management is watching the gap widen with each cycle.

For born-agentic businesses, the compounding starts earlier and at a steeper angle, because the first phase of the maturity journey: the phase that Chapter 4 identified as delivering 45 percent profitability uplift, is not separated from full integration by a migration journey. It is the starting point.[106] The born-agentic firm begins closer to the second phase of the J-curve where the 111 percent uplift lives, with no legacy transition costs reducing the net return on its investment.[107]

THE PARTNER IMPLICATION: TWO DIFFERENT CLIENTS, TWO DIFFERENT RELATIONSHIPS

The convergence of the adoption gap and the emergence of the born-agentic cohort have a direct implication for how technology partners should segment their client base and structure their service offerings. The legacy SMB engaged in an AI migration journey is a fundamentally different client from the born-agentic startup, and the two require different service architectures.

The legacy client needs a migration partner capable of:

- Assessing existing infrastructure
- Sequencing the integration of AI capabilities in a way that minimizes disruption
- Managing the change management requirements of workflow redesign with existing staff
- Building the data foundations that AI systems require from a starting point of fragmented legacy records

This is complex, high-value work with strong switching costs and long engagement timelines; the organizational architecture work that Chapter 4 identified as the second phase of the J-curve.

The born-agentic client needs an architecture partner from day one, and this is where Forrester's prediction becomes directly relevant. Forrester finds that **three out of four firms attempting to build agentic architectures on their own will fail,** given the stack complexity involved: multiple AI models, retrieval-augmented generation infrastructure, data architecture decisions, governance frameworks and niche expertise requirements that no small business team can reasonably develop internally from inception.[108] The born-agentic business that partners with an MIP from its first day does not just get better tools configured faster. It gets the governance architecture, the data foundations and the agent orchestration framework that turns AI capability into compounding competitive advantage; and it gets them before its competitors have even begun their migration journeys.

The providers who recognize this distinction early, who build differentiated service packages for born-agentic startups distinct from their legacy SMB migration services, will be positioned to capture the full range of the opportunity that the born-agentic shift is creating. These clients have lower implementation complexity but higher strategic advisory requirements, enabling premium pricing with lower delivery costs and a partner relationship embedded in the institutional knowledge of how the business was designed to operate from its very first day.

Chapter 7 turns to the risk dimension of this same shift: the finding that every AI agent, every automated workflow and every born-agentic architecture also expands the attack surface in direct proportion to its productivity value; and that the cybersecurity implications of digital labor adoption are not a separate concern to be managed alongside the AI strategy, but an inseparable dimension of the investment itself.

105 Salesforce. "New Research Reveals SMBs with AI Adoption See Stronger Revenue Growth." Salesforce, 2025.

106 O'Mahony, John, et al. "The AI Edge for Small Business: Increased SMB AI Adoption Can Add $44 Billion to Australia's Economy." Deloitte Access Economics, 25 Nov. 2025.

107 Alegbeh, Alchad, and Marvin Cruz. "Digital Transformation: How Small Businesses in Canada Are Leveraging AI and Technology for Growth and Productivity." Canadian Federation of Independent Business, 29 Sept. 2025.

108 Forrester Research. "Predictions 2025: An AI Reality Check Paves the Path for Success." Forrester, 2025.

Part 03.
The Risk Equation

Agentic labor does not arrive in the small business economy as a pure productivity benefit. It arrives as a dual transformation.

The first two parts of this report have examined the economic case for agentic labor: the productivity gains, the valuation implications, the compounding returns available to businesses that complete the integration journey, and the structural advantage available to those who build on AI-native architectures from the start. Part Three examines the other side of that same investment.

Agentic labor does not arrive in the small business economy as a pure productivity benefit. It arrives as a dual transformation: simultaneously expanding what a business can produce and expanding the surface across which it can be attacked, compromised or destabilized. Every AI agent deployed, every automated workflow activated, every born-agentic architecture built adds capability on one axis and vulnerability on another; and in many businesses currently deploying this digital labor, those two axes are being managed by different people, on different timelines, with different levels of urgency.

Part Three examines three dimensions of the risk equation that the productivity narrative consistently underweights: the cybersecurity exposure that grows in direct proportion to the value of agentic labor adoption; the data readiness deficit that undermines AI performance while simultaneously creating new compliance vulnerabilities; and the governance gap that represents both the greatest threat to agentic AI deployments and the greatest opportunity for technology partners willing to fill it.

Chapter 07.

Productivity and Vulnerability Are the Same Investment

There is a version of the agentic labor adoption story told almost exclusively in the language of gain. Productivity lifts, profitability uplifts, time dividends, valuation multiples, competitive advantage; the economic case is real, and the preceding chapters have made it clear with data. But every one of those gains is predicated on a digital infrastructure that is expanding in complexity, connectivity and autonomous capability at a pace that the security and governance architectures surrounding it are not matching. The businesses that understand this are building a more durable competitive advantage, because they are treating the risk dimension of digital labor adoption as inseparable from the productivity dimension rather than as a separate problem to be managed later.

The ones that do not understand this are carrying a compounding liability they cannot currently see; one that becomes visible only when a single incident wipes out months of carefully accumulated productivity gains in a matter of hours.

THE ATTACK SURFACE EXPANDS WITH THE OPPORTUNITY

The foundational dynamic of this chapter deserves to be stated plainly: every AI agent, every automated workflow and every piece of digital infrastructure that makes an SMB more productive also makes it a more attractive and accessible target. This is not a speculative future risk. It is an accelerating present condition.

CrowdStrike's 2026 Global Threat Report puts numbers on the acceleration that are difficult to look away from. AI-powered cyberattacks surged 89 percent year-over-year. Average breakout times: the interval between initial system compromise and lateral movement across a network dropped to 29 minutes in 2025, 65 percent faster than in 2024, with some attacks unfolding in mere seconds.[109] Attackers are deploying the same AI capabilities that make businesses more productive to make their intrusions faster, more targeted and harder to detect. State-sponsored actors including Russia's Fancy Bear are leveraging large language model-enabled malware for automation and North Korea's Famous Chollima has escalated insider operations using AI-generated personas.[110] The tactical arms race between AI-enabled defense and AI-enabled offense is the current operational environment for every business running connected digital infrastructure.

For small and medium businesses, the exposure is disproportionate in ways the aggregate statistics frequently obscure. Halcyon's research on SMB cybersecurity finds that 88 percent of SMB breaches now involve ransomware; a rate that far exceeds the 39 percent seen in larger enterprises, reflecting the reality that smaller businesses typically carry fewer dedicated security resources, less mature detection and response capabilities, and a lower organizational threshold for the disruption that a successful attack delivers.[111]

The gap reflects a structural reality: smaller businesses carry fewer dedicated security resources, less mature detection and response capabilities, and a lower organizational threshold for the disruption a successful attack delivers. Attackers know this, and they price their targeting decisions accordingly. The financial consequences compound that vulnerability. Sophos's 2025 State of Ransomware report, drawing on responses from 3,000 IT and security professionals across the globe, documents that the average cost to recover from a ransomware attack approximated $1.53 million in 2025, and for SMBs without enterprise-scale backup infrastructure and incident response capacity, even a fraction of that figure can represent an existential financial event.[112] ConnectWise's global SMB cybersecurity research adds a further dimension to the exposure picture: 47 percent of small businesses were hit by ransomware in the past year alone, and the average ransom payment increased 500 percent between 2020 and 2024 to reach $2 million.[113]

Smaller businesses carry fewer dedicated security resources, less mature detection and response capabilities, and a lower organizational threshold for the disruption a successful attack delivers.

The ASUS Future of SMB Report 2025 makes the exposure personal. 27 percent of SMB leaders report feeling unsafe about their current cybersecurity posture. 28 percent have already experienced a cyberattack.[114] These are the same businesses deploying AI tools, adopting automated workflows and in the case of born-agentic organizations, building entire operational architectures on digital foundations.

The productivity investment and the security exposure are not separate decisions. They are the same decision, encountered from different angles, and the businesses treating them as separate are managing one half of the equation while the other compounds against them.

109 CrowdStrike. "2026 Global Threat Report." CrowdStrike, 2026.

110 CrowdStrike. "2026 Global Threat Report." CrowdStrike, 2026.

111 Halcyon. "Small and Medium Businesses Under Siege." Halcyon, 2025.

112 Sophos. "State of Ransomware 2025." Sophos, 2025.

113 ConnectWise. "SMB Cybersecurity Statistics and Trends." ConnectWise, 2025.

114 ASUS Business. "The Future of SMB Report 2025: Harnessing the Potential of AI PCs." ASUS, 2025.

AGENTIC AI CHANGES THE RISK PROFILE FUNDAMENTALLY

If standard digital tool adoption expands the attack surface, agentic AI expands it by an order of magnitude, because agents, unlike passive software tools, take actions. They make purchases. They send communications. They access files, modify records, authenticate to third-party services and execute multi-step workflows with limited or no human intervention at each step.

Each of those capabilities is a genuine productivity multiplier. Each is also a new vector through which a compromised system, a malicious prompt injection or a governance failure can produce real-world consequences that a traditional software vulnerability cannot.

Nathan Baschez, founder of an AI-powered word processor company, describes a "Lethal Trifecta" framework from software developer Simon Willison: an agent becomes dangerous when it simultaneously has access to private data, processes untrusted external content and can communicate outward. Removing any one of those three conditions significantly drops the overall risk profile; a finding that has direct implications for how any AI agent deployment should be designed and governed.[67] Even a leading AI agent user like Baschez, who reports moving 15 to 20 times faster with AI assistance, maintains explicit security rules ensuring agents never hold all three risk conditions simultaneously.[115]

Gartner has formalized the governance gap at the enterprise level with a prediction that should concentrate the attention of every technology partner working with SMB clients: more than 40 percent of AI-related data breaches by 2027 will arise from improper cross-border use of generative AI, driven by end-user adoption outpacing the policies, controls and monitoring systems required to manage it safely.[116] A separate analysis finds that through 2026, more than 40 percent of agentic AI projects will be abandoned before completion, with the primary causes being governance gaps, data infrastructure inadequacies and the structural mismatch between deployment ambition and operational readiness.[117] These predictions are about organizational failure; the same last mile problem that Chapter 3 identified as the barrier to AI value realization, now carrying security consequences rather than just productivity ones.

Haider Pasha, EMEA CISO at Palo Alto Networks, has described the governance challenge of agentic AI in terms that should frame every technology partner's client conversation. Securing an AI agent, he argues, is analogous to managing a new employee who has been given broad system access before their authorization boundaries have been established. The risks: memory misuse, objective drift, prompt injection vulnerabilities are the predictable consequences of deploying systems with tool access and autonomous decision-making capability before the organizational governance frameworks required to contain them are in place.

THE TRUST GAP NOBODY IS CLOSING FAST ENOUGH

The scale of the trust deficit surrounding agentic AI provides perhaps the most revealing window into how far the governance infrastructure of most organizations lags behind their deployment ambitions. The Harvard Business Review Analytic Services study of more than 600 technology decision-makers finds that only 6 percent of organizations fully trust AI agents to autonomously handle core business processes.[118] Yet 86 percent of those same organizations plan to increase their investment in agentic AI over the next two years, and 74 percent are already working on or planning enterprise orchestration frameworks to prepare for broader agentic deployment.[119]

> Organizations are deploying capability they do not yet fully trust.

That gap — 6 percent trusting, 86 percent investing — is a description of the governance deficit that is generating the breach statistics and project failure rates Gartner is forecasting. Organizations are deploying capability they do not yet fully trust, at a pace that their risk management practices cannot match. For SMBs, which lack the internal resources to build governance infrastructure at the same speed as their tool adoption, the gap is wider still. The U.S. Chamber of Commerce's technology survey found that 58 percent of SMBs have already adopted AI in some form, up from 40 percent just one year earlier.[120] The pace of that adoption has outrun the pace of policy development in many of those businesses.

ConnectWise's cybersecurity data finds that 49 percent of SMBs do not have AI-specific security policies in place.[121] That gap: 58 percent adopting, 49 percent ungoverned, is the attack surface of the next wave of SMB cyber incidents, already built and largely undefended.

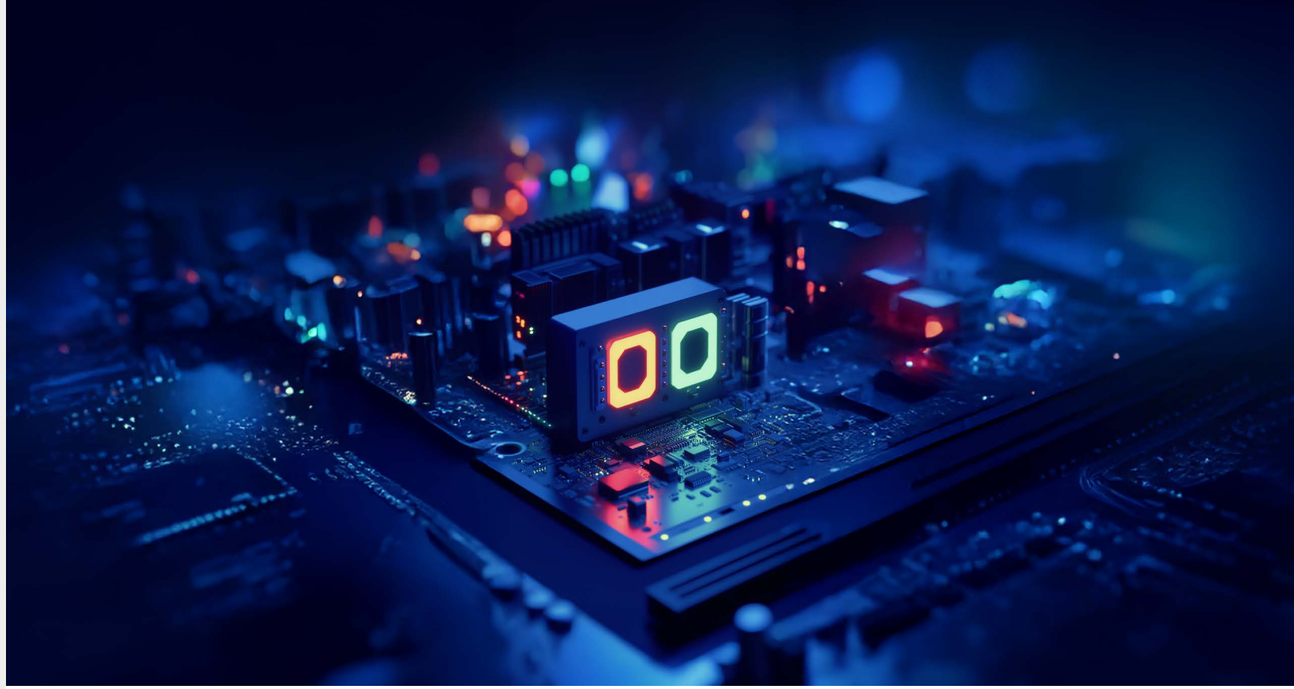

THE SHADOW AI PROBLEM ALREADY INSIDE

One of the most underappreciated dimensions of the SMB security exposure is not the AI that businesses are deliberately deploying but the AI they are not tracking. Gartner predicts that by 2030, more than 40 percent of global organizations will experience security and compliance incidents driven by shadow AI: the use of AI tools by employees outside of approved IT oversight, without governance policies, data controls or audit trails.[122] Among organizations surveyed today, 69 percent have evidence or strong suspicion that employees are already using unauthorized AI tools at work.[123]

For SMBs, the shadow AI problem is acute precisely because the governance infrastructure required to detect and manage it is the same infrastructure that most small businesses have not yet built. A customer service employee who has discovered that a consumer AI tool can draft responses faster than their approved workflow is not acting maliciously. They are solving a productivity problem with the most accessible available tool. But in doing so, they are introducing company data — customer records, transaction histories, internal communications — into AI systems the business has not evaluated, approved or configured with the data handling requirements that compliance and security demand. Microsoft's own research finds that 71 percent of UK workers admit using unauthorized AI tools at work, with 22 percent employing them for critical financial tasks.[124] The shadow AI population is not small, and it is not confined to any single region or industry.

115 Griffin, Anne. "4 Businesses Adopting AI Agents to Transform How They Operate." Forbes, 25 Feb. 2026.

116 Gartner. "Gartner Predicts 40% of AI Data Breaches Will Arise from Cross-Border GenAI Misuse by 2027." Gartner, 17 Feb. 2025.

117 Gartner. "Gartner Identifies the Top 10 Strategic Technology Trends for 2025." Gartner, 21 Oct. 2024.

118 Harvard Business Review Analytic Services. "From the Edge to the Core: Bringing Agentic AI to the Heart of the Enterprise." Harvard Business Review Analytic Services, Dec. 2025.

119 Harvard Business Review Analytic Services. "From the Edge to the Core: Bringing Agentic AI to the Heart of the Enterprise." Harvard Business Review Analytic Services, Dec. 2025.

120 U.S. Chamber of Commerce. "Empowering Small Business: The Impact of Technology on U.S. Small Business." U.S. Chamber of Commerce, 13 Aug. 2025.

121 ConnectWise. "SMB Cybersecurity Statistics and Trends." ConnectWise, 2025.

122 Gartner. "Gartner Predicts 40% of AI Data Breaches Will Arise from Cross-Border GenAI Misuse by 2027." Gartner, 17 Feb. 2025.

123 Gartner. "Gartner Predicts 40% of AI Data Breaches Will Arise from Cross-Border GenAI Misuse by 2027." Gartner, 17 Feb. 2025.

124 Microsoft. "Rise in 'Shadow AI' Tools Raising Security Concerns for UK." Microsoft UK Stories, Oct. 2025.

THE COMPOUNDING DANGER OF DEFERRAL

The risk calculus of digital labor adoption is frequently framed as a binary: adopt and face the security implications or defer and avoid them. The research makes clear this framing is structurally wrong. Deferring digital labor adoption does not eliminate exposure to the security risks of the digital labor era. It eliminates the competitive benefit while leaving the security exposure substantially intact, because the threat environment is not a function of a given business's own AI deployment. It is a function of the broader digitization of the economy, which exposes every connected business regardless of whether it is actively deploying AI.

Bain's recent Technology Report frames the compounding dynamic on the competitive side with directness that has no real counterargument: AI leaders have moved from pilots to profits, achieving 10 to 25 percent EBITDA gains by scaling AI across core workflows; and companies still in the piloting phase are, in Bain's characterization, dangerously behind.[125] The businesses that deferred adoption to manage risk did not necessarily reduce their risk exposure. They reduced their returns while the gap between themselves and businesses that made the investment continued to widen each quarter.

MIT's Project NANDA research provides the granular failure analysis that explains why most deferrals do not produce the safety they are intended to create. Its 2025 GenAI Divide report found that 95 percent of organizations deploying generative AI are generating zero measurable return on investment, while 5 percent of integrated deployments are extracting

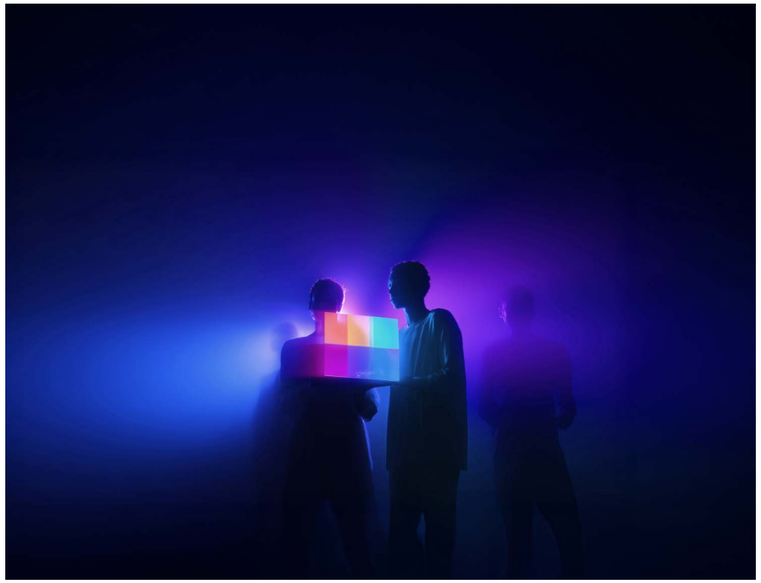

what the research describes as millions in value.[126] Critically for technology partners, the research identifies external partnerships as consistently outperforming purely internal AI builds — positioning the managed intelligence relationship not as a luxury but as the primary mechanism through which value capture becomes achievable and security governance becomes sustainable.[126] The organizations in the successful 5 percent are not simply more profitable. They are also more secure, because the organizational discipline required to achieve deep AI integration is the same discipline required to govern it safely.

THE UNIFIED RESPONSE: INTEGRATING SECURITY INTO THE DIGITAL LABOR STACK

The implication of the research across all these dimensions is not that SMBs should slow their digital labor adoption out of security caution. The businesses that deferred cloud adoption for security reasons in 2012 are not more secure today; they are simply behind. The implication is that the businesses deploying digital labor and the technology partners helping them do so need to treat security governance not as a separate work stream to be addressed after deployment but as a foundational dimension of the deployment itself.

Deloitte's State of AI in the Enterprise 2026 documents the governance gap with the same precision it brings to the productivity opportunity: close to three-quarters of companies plan to deploy agentic AI within two years, but only 21 percent currently report a mature model for governing autonomous agents.[127] The businesses in that 21 percent are not simply better protected against the breach scenarios that Gartner is forecasting. They are the businesses able to deploy AI agents with broad autonomous capability that generates the returns Bain documents, because they have built the governance architecture that makes that autonomy safe to extend at scale.

For MSPs and MIPs, the security dimension of agentic labor adoption is not a separate service line to be sold alongside the productivity offering.

For MSPs and MIPs, the security dimension of agentic labor adoption is not a separate service line to be sold alongside the productivity offering.

It is the moat. Every AI agent deployed, every automated workflow activated, every born-agentic architecture built requires someone to govern it; to define its boundaries, monitor its behavior, manage its access credentials, audit its actions and maintain the policy frameworks that keep it aligned with the business's compliance obligations and risk appetite. That work cannot be done by the SMB alone. It requires the kind of continuous, expert-driven oversight that is the defining characteristic of a managed intelligence relationship. The providers who build that capability and bundle it inseparably with their agentic labor deployment practice are the only providers offering what the SMB market demonstrably needs and has no other reliable source for.

Chapter 8 explores a closely related dimension of this same problem: the finding that SMBs want AI's outputs without having built the data foundations that AI requires, and that the gap between what they expect digital labor to deliver and what their actual data infrastructure currently enables is the single most common reason digital labor investments underdeliver.

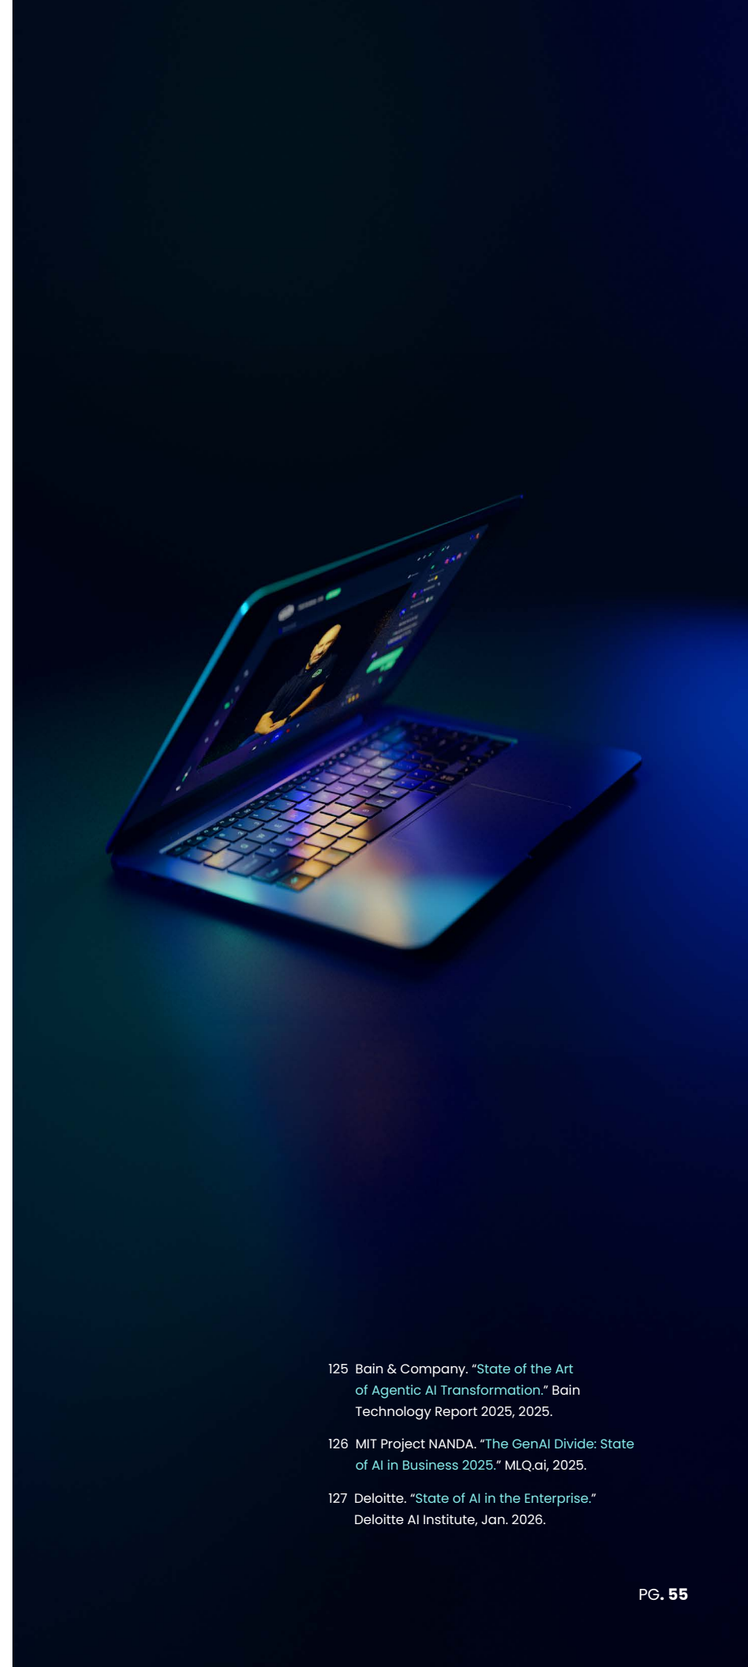

125 Bain & Company. "State of the Art of Agentic AI Transformation." Bain Technology Report 2025, 2025.

126 MIT Project NANDA. "The GenAI Divide: State of AI in Business 2025." MLQ.ai, 2025.

127 Deloitte. "State of AI in the Enterprise." Deloitte AI Institute, Jan. 2026.

Chapter 08.

The Data Readiness Chasm

The answer, in the majority of SMB environments, is data.

There is a question that cuts to the heart of why most AI investments in the small business economy underdeliver, and it is not a question about technology. It is not about which AI model a business chooses, which software vendor they partner with or which workflows they automate first. It is a question about infrastructure that most businesses do not think to ask until they are several months into a deployment that is producing underwhelming results: what data is the AI working with?

The answer, in the majority of SMB environments, is data that is fragmented across disconnected systems, inconsistent in format and quality, siloed by department or function and accessible only through manual extraction processes that were never designed for the kind of real-time, cross-system intelligence that AI agents require to operate effectively. Agentic labor does not make fragmented data infrastructure less of a problem. It makes it more of one, because every AI agent deployed into a broken data environment inherits the limitations of that environment, and then executes on them at machine speed.

The data readiness chasm is not the most visible barrier to AI value realization in the SMB economy. It is not the barrier that generates the most attention in vendor marketing, the most discussion at industry conferences or the most anxiety in small business owner surveys. But it is consistently the most common proximate cause of AI pilots that stall, deployments that underperform and ROI commitments that go unfulfilled. Understanding it, and knowing how to close it, is the difference between the technology partner who gets called when an AI initiative is struggling and the one who was there at the beginning, building the foundations that made the initiative work.

WHAT HIGH PERFORMERS DO DIFFERENTLY

The most important data point in McKinsey's 2025 State of AI global survey is what separates the approximately six percent of organizations generating measurable enterprise-wide financial impact from the 94 percent that are not.[128] McKinsey calls this cohort AI high performers, and their distinguishing characteristics have been remarkably consistent across the firm's annual surveys: they treat AI as a vehicle for genuine workflow transformation rather than incremental efficiency; they have senior leadership actively and visibly championing AI initiatives; they invest meaningfully: more than one-third commit over 20 percent of their digital budgets to AI; and they are nearly three times more likely than their peers to have fundamentally redesigned the workflows through which AI operates.[129]

Each of those characteristics matters. But running underneath all of them is a precondition that the productivity narrative consistently underweights high performers have built the data infrastructure that makes those characteristics actionable. Redesigning a workflow requires understanding what data flows through it and where that data lives. Scaling AI across functions requires those functions to share accessible, governed data. Deploying AI agents that can act autonomously requires data environments clean enough and integrated enough that agent decisions are grounded in accurate, current information rather than stale, incomplete records.

Salesforce's global SMB Trends research documents the behavioral signature of this precondition at the small business level. Among growing SMBs, the cohort already capturing AI returns, 74 percent are actively increasing their investment in data management. Among declining SMBs, that figure drops to 47 percent.[130] The gap between those two numbers is a foundational infrastructure gap, and it is compounding in the same way the productivity gap compounds: every quarter that a growing SMB invests in its data foundation while a declining competitor defers that investment, the distance between their AI capabilities widens.

THE ASPIRATION-REALITY DIVIDE

The data readiness problem is made structurally more complex by the gap between what SMBs say they want AI to do and where their actual digital infrastructure currently operates. Survey after survey captures SMB leaders expressing high aspirations for AI in customer experience, sales forecasting, marketing personalization and strategic decision support; all functions that require clean, integrated, real-time data to deliver on their promise. The actual digital tool adoption picture tells a different story.

Additional research on real SMB digital tool usage finds that the operational functions where small businesses have actually achieved meaningful digital adoption are considerably more basic: payment collection (41 percent of SMBs have digitized this function to a meaningful degree), billing and invoicing (46 percent), market automation (36 percent) and scheduling (33 percent).[131] These are not the aspirational AI use cases that appear in SMB survey responses. They are the foundational transactional functions where data is cleanest, processes are most defined and business outcomes are most measurable; and they represent the realistic beachhead from which data-ready AI adoption can begin.

The implication is that most SMBs are not starting from the data foundation that their AI aspirations require, but a patchwork of digitized transactional functions, surrounded by manual processes, disconnected systems and data that exists in formats their AI tools cannot reliably consume. The SMB that tells a technology partner it wants AI to improve strategic decision-making is often, in practice, operating with data systems that struggle to answer basic operational questions consistently. Closing that gap is the work that precedes every AI deployment that delivers returns; and it is work that most SMBs have neither the internal expertise nor the organizational capacity to do on their own.

128 Singla, Alex, et al. "The State of AI in 2025: Agents, Innovation, and Transformation." McKinsey & Company, Nov. 2025.

129 Singla, Alex, et al. "The State of AI in 2025: Agents, Innovation, and Transformation." McKinsey & Company, Nov. 2025.

130 Salesforce. "Small and Medium Business Trends Report, 6th Edition." Salesforce, 2025.

131 vcita. "inTandem SMB Digital Adoption Report." vcita, 2025.

The data readiness problem that undermines basic AI adoption becomes existentially more consequential in an agentic context.

THE CONFIDENCE BARRIER THAT CONCEALS THE PROBLEM

What makes the data readiness chasm particularly difficult to address is that most SMB leaders do not know they have a data problem until an AI deployment makes it visible. The Salesforce research is striking on this point: 84 percent of SMB leaders agree that clean, complete data is critical to business success, and 82 percent believe improving data quality would directly boost their operational efficiency.[132] These are not the responses of leaders who are dismissive of data quality. Yet the same survey finds that only 47 percent of SMBs say they currently make decisions fully based on data; meaning more than half are operating on intuition, partial information or data they do not fully trust.[133]

That gap: 84 percent believing clean data is critical, 47 percent operating on clean data, is the confidence illusion that allows AI investments to be made without the data investments that would make them productive. The SMB that believes its data is good enough because nothing has broken yet is the SMB that discovers what "not good enough" means when an AI agent begins making autonomous decisions based on it. A customer service agent that references outdated account information. An inventory agent that reorders based on stale demand signals. A scheduling system that operates on contact records that have not been updated in two years. Each of these is a data readiness failure that the AI has made immediately and expensively visible.

MIT's Project NANDA research identifies this dynamic as the defining characteristic of the 95 percent of AI deployments that generate zero measurable financial return.[134] The report, based on analysis of 300 enterprise AI deployments, is explicit: the failures are not caused by inadequate AI models. They are caused by inadequate integration of AI into existing workflows, and inadequate workflows are almost always downstream consequences of inadequate data infrastructure.[135] The businesses that are generating returns are not operating better AI; they are operating on better data, with better-defined processes and with governance frameworks that make those processes legible to AI systems. The technology is largely the same. The foundation is entirely different.

THE AGENTIC STAKES

The data readiness problem that undermines basic AI adoption becomes existentially more consequential in an agentic context. A generative AI tool that produces a mediocre output because it was working with poor data is a productivity disappointment. An AI agent that takes an incorrect autonomous action because it was working with poor data is a business liability; and potentially a security incident, a compliance violation or a customer relationship crisis, depending on what the agent was authorized to do.

The Business Insider analysis of agentic AI enterprise data requirements makes this structural dependency explicit: AI agents operating in fast-moving business environments require dynamic, real-time data rather than periodically updated static datasets; orchestrated rather than siloed data access across the systems the agent must navigate; and proactive data governance rather than reactive problem management.[136] None of these requirements describes the data environment that exists in most SMBs today. Most small businesses operate with data that is updated periodically rather than in real time, that lives in separate systems that do not communicate with each other reliably and that is governed, if at all, by informal conventions rather than systematic policies.

Deloitte's report captures the governance dimension: close to three-quarters of organizations plan to deploy agentic AI within two years, yet only 21 percent report having a mature model for governing autonomous agents.[137] The 79 percent that lack mature governance are missing the data infrastructure; the audit trails, the access controls, the data quality monitoring and the integration architecture that governance frameworks require to function. Building the governance and building the data foundation are the same work, approached from different angles.

THE BILLABLE OPPORTUNITY THAT MOST PARTNERS ARE MISSING

The data readiness chasm is not a problem to be solved before the AI conversation begins. It is the AI conversation, and for technology partners willing to lead with it rather than skip past it, it represents the most durable billable engagement category in the current market.

IDC's research documents the demand signal for this work: 70 percent of SMBs will demand clear AI use cases from vendors and managed service providers before investing in new AI or generative AI technologies.[138] The implied question in that demand is not simply "what can AI do?" It is "what can AI do for my specific business, with my specific data, in my specific operational context?" Answering that question honestly almost always begins with a data readiness assessment: an audit of what data the business has, where it lives, how clean it is, whether systems can communicate with each other, and what the integration work required to make AI agents effective would entail.

That assessment is not glamorous work. It does not involve demonstrating impressive AI capabilities or showcasing cutting-edge technology. It involves the unglamorous but essential work of data hygiene, identity integration, API connectivity and system architecture that determines whether every subsequent AI investment delivers or disappoints. McKinsey's research is unambiguous on this point: high performers are nearly three times more likely to have fundamentally redesigned their workflows; and workflow redesign, done properly, begins with understanding and rationalizing the data flows those workflows depend on.[139]

For MSPs and MIPs, the strategic positioning implication is plain. The technology partner who leads every new AI engagement with a data and integration readiness assessment, packaged as a fixed-fee discovery offering that maps the client's data quality, system connectivity and governance gaps, is not a delay in the AI conversation. It is leading with one with one of its most critical foundational elements. That assessment creates a natural two-phase engagement architecture: data foundation work that generates immediate, measurable improvement in operational reliability, followed by AI agent deployment into an environment that is ready to support it. The partners who build this into their standard engagement model will build the kind of embedded, institutional understanding of each client's data environment that makes switching to a different provider costly and unlikely.

Chapter 9 examines how the cybersecurity dimension of this same work creates an additional strategic lever; the finding that security is not merely a parallel service line to AI adoption but the fastest trust-building path to it, and that the providers who understand this have a structural head start over those still treating security and AI as separate conversations.

132 Salesforce. "Small and Medium Business Trends Report, 6th Edition." Salesforce, 2025.

133 Salesforce. "Small and Medium Business Trends Report, 6th Edition." Salesforce, 2025.

134 MIT Project NANDA. "The GenAI Divide: State of AI in Business 2025." MLQ.ai, 2025.

135 MIT Project NANDA. "The GenAI Divide: State of AI in Business 2025." MLQ.ai, 2025.

136 Business Insider. "Three Critical Data Shifts Your Company Needs to Thrive in the Agentic AI Era." Business Insider, 2025.

137 Rowan, Jim, et al. "State of AI in the Enterprise." Deloitte AI Institute, Jan. 2026.

138 Evans, Katie, et al. "IDC FutureScape: Worldwide Small and Medium-Sized Business 2025 Predictions." IDC, 2024.

139 Singla, Alex, et al. "The State of AI in 2025: Agents, Innovation, and Transformation." McKinsey & Company, Nov. 2025.

Chapter 09.

The Cybersecurity Trojan Horse

For the MSPs and MIPs who understand it, it is the fastest available path from security provider to agentic labor partner.

There is an argument about AI adoption that almost never appears in industry research, vendor marketing or SMB owner surveys, and yet it is hiding in plain sight in almost every managed services engagement. It is an argument about trust; specifically, about where trust in autonomous AI decision-making already exists in the small business economy, and what that existing trust unlocks for the providers who recognize and build on it.

The argument is this: many SMBs are already running autonomous AI agents that make real-time decisions without human approval, in a domain where the stakes of those decisions are immediate and consequential. Those agents are not productivity tools or workflow automation platforms. They are cybersecurity systems: threat detection engines, endpoint protection platforms, incident response tools; and they operate continuously, autonomously, and in the background of nearly every connected business in today's economy. The SMB owner who tells a technology partner that they are nervous about AI agents making autonomous decisions in their business has, in most cases, been relying on autonomous AI agents to protect their business for years. They simply have not framed it that way.

That framing gap is the Trojan Horse. And for the MSPs and MIPs who understand it, it is the fastest available path from security provider to agentic labor partner, without a cold start, without a trust-building journey from scratch and without competing on a dimension where most of the market is already crowded.

THE TRUST THAT ALREADY EXISTS

The endpoint protection platform monitoring a client's devices at 2 AM on a Sunday is not waiting for a human to approve its decision to quarantine a suspicious process. The intrusion detection system scanning network traffic for anomalous behavior is not pausing its analysis to ask whether it should escalate a potential breach. The email security layer that intercepts a phishing attempt before it reaches an employee's inbox made that decision autonomously, at machine speed, before any human was aware the attempt had occurred. These systems plan, assess, and act, which is precisely how Gartner defines the category of agentic AI.[140]

SMBs that have adopted managed security services have already approved the governance model that underlies all agentic AI deployment: a trusted external partner governs a system that takes autonomous actions within defined parameters, with human oversight reserved for escalation and exception rather than routine operation. The audit trails, the alerting thresholds, the response playbooks, the containment procedures; these are the governance architecture of an autonomous agent. They are already built, already trusted and already operational in the security domain. The only question is whether the technology partner helping to govern them has recognized that they are also the template for every other autonomous agent deployment the client will ever make.

THE TRUST DEFICIT THAT CREATES THE OPPORTUNITY

If the trust that makes this argument powerful is the trust that already exists, the autonomous AI in security that SMBs have already accepted, then the trust that creates the market opportunity is the trust that does not yet exist: the 73 percent of SMBs that are not fully confident their MSP could defend them against a cyberattack, the 47 percent that would switch providers for a stronger cybersecurity solution, and the 32 percent that would hold their MSP solely responsible in the event of a breach, with 79 percent of those open to legal action.[141]

Read together, those numbers describe a market that is simultaneously dependent on its MSPs for security and deeply skeptical about whether those MSPs can deliver it. That is not a comfortable position for the average managed services provider. But it is a structural opening for providers who can close that confidence gap, because the SMB that decides to switch providers for stronger security is looking for a partner relationship that they can take with them into every subsequent technology decision they make, including the digital labor investments that represent the largest technology spending opportunity of the next decade.

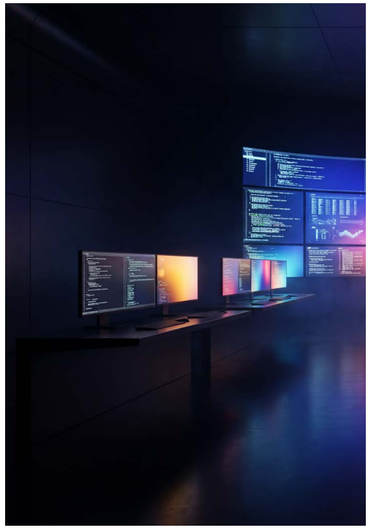

ConnectWise's 2025 State of SMB Cybersecurity report adds a further dimension that sharpens this dynamic: 83 percent of SMBs believe that AI and generative AI have increased their organization's threat exposure.[142] That figure is the manifestation of a present-tense anxiety held by more than four in five small businesses, and it is not being adequately addressed. Despite that near-universal recognition of elevated AI-driven threat exposure, only 51 percent of SMBs have implemented security policies or practices specifically governing AI and generative AI use in their organizations.[143] The gap between 83 percent recognizing the risk and 51 percent having governance around it is an execution gap, and execution gaps in security are precisely what managed services providers exist to close.

140 Gartner. "Gartner Identifies the Top 10 Strategic Technology Trends for 2025." Gartner, 21 Oct. 2024.

141 ConnectWise. "SMB Cybersecurity Statistics and Trends." ConnectWise, 2025.

142 ConnectWise. "SMB Cybersecurity Statistics and Trends." ConnectWise, 2025.

143 ConnectWise. "SMB Cybersecurity Statistics and Trends." ConnectWise, 2025.

POLICY IS THE ENTRY POINT
The 49 percent of SMBs without AI-specific security policies represent an immediately addressable, high-urgency service opportunity for every partner who can articulate why those policies matter and what they need to contain. This is operational risk management for businesses that are, by their own admission, deploying AI tools; shadow AI and sanctioned alike; without the governance frameworks required to manage the exposure those tools create.

As Chapter 7 documented, the shadow AI population in the SMB economy is substantial and growing. The Microsoft research finding that 71 percent of workers use unauthorized AI tools at work, with 22 percent doing so for financial tasks is a behavioral pattern that follows wherever consumer AI tools are accessible and organizational policies have not kept pace with adoption speed.[144] Every SMB in that 49 percent without AI security policies has employees making those decisions unilaterally, routing company data through systems the business has not evaluated, approved or configured for its compliance requirements. The MSP or MIP that helps an SMB build its first AI security policy framework is the organization that establishes themselves as the authoritative voice on what the client's AI environment looks like, what is permitted within it and how it will be governed going forward. That is the foundation of the managed intelligence relationship.

Additional research adds important texture to where this exposure is concentrated. The data finds that 84 percent of SMB leaders feel safe or neutral about their cybersecurity posture, while simultaneously revealing that 16 percent have already experienced a cyberattack.[145] That confidence gap has real consequences: Microsoft's own SMB cybersecurity research finds that 31 percent of SMBs have been victims of a cyberattack; including ransomware, phishing or data breaches; underscoring how significantly small businesses underestimate their exposure.[146]

Building organizational trust in autonomous AI decision-making is the hardest part.

The confidence and the reality are sharply misaligned. The SMB that feels safe because it has not been breached yet is the SMB that will be most shaken when a breach occurs and most receptive to a trusted partner who can explain what happened, why their existing posture was insufficient and what a more defensible architecture would look like; one that explicitly addresses the AI-driven threat vectors that are now responsible for 89 percent of year-over-year growth in cyberattacks.[147]

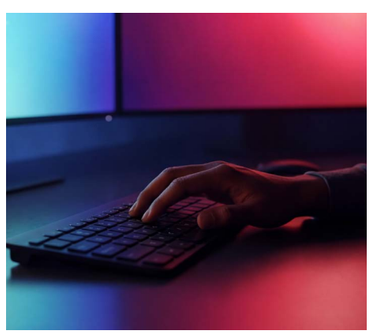

THE GOVERNANCE MODEL IS ALREADY PROVEN
The deepest strategic insight of this chapter reveals what the security conversation unlocks for providers who frame it correctly. The governance model that a technology partner builds around a client's cybersecurity environment: the monitoring, the alerting, the incident response playbooks, the access controls, the audit trails, the human escalation protocols; is structurally identical to the governance model required for any other autonomous agent deployment. It is the same model applied to a different domain.

When an MSP or MIP monitors endpoint behavior and escalates anomalies to a human security analyst, it is performing the same governance function that will be required when an AI agent is managing invoice processing and needs oversight for exception cases. When a partner maintains audit trails of security events across a client environment, it is building the same institutional capability required to audit agent actions in billing, customer service or scheduling workflows. When they define the parameters within which an automated threat response tool can act autonomously and the thresholds at which it must escalate for human decision-making, it is designing the same bounded autonomy architecture that responsible agentic AI deployment requires in every operational context.

The providers who frame their existing security practice this way are making a strategic argument with implications that extend well beyond security. The trust model governing a client's cybersecurity agents, including the access controls, escalation policies, human oversight checkpoints and audit trails, is the precise blueprint for governing every other autonomous agent they deploy. Building organizational trust in autonomous AI decision-making is the hardest part of the digital labor adoption journey. The client who has already built that trust in a security context does not need to build it again from scratch for billing automation or customer service workflows. They need to recognize that it already exists, extend it to new domains with the same governance discipline and do so with the partner who has already demonstrated the capability to manage it.

THE STRUCTURAL HEAD START
For MSPs and MIPs that have been building their practices around cybersecurity, the digital labor era represents not a disruption of their existing value proposition but its elevation. The threat detection, the policy governance, the compliance management, the incident response, the security architecture work; all of it is upstream infrastructure for the autonomous agent economy that is arriving whether SMBs are ready for it or not. Forrester's finding that three out of four organizations attempting to build agentic architectures on their own will fail, given the stack complexity involved, applies with equal force to security governance.[148] Building a defensible AI governance framework from scratch is not a project that a 15-person professional services firm or a 40-person manufacturing operation can execute with internal resources. It requires external expertise, and the providers who have already built that expertise in the security domain have the most credible claim to it in the digital labor domain as well.

The providers who do not recognize this; who continue to treat cybersecurity as a separate practice line from AI adoption, selling security to one stakeholder and AI productivity to another, never connecting the governance conversation to the deployment conversation; will cede that integrated relationship to competitors who do. The Trojan Horse works in both directions. For the providers who enter the agentic labor conversation through the security door, it delivers the client relationship intact, with trust already established. For the providers who wait for the AI productivity conversation to arrive on its own terms, the security-first competitor will already be inside.

Part Four turns from the risk environment facing SMBs to the supply-side reckoning facing the providers serving them; the finding that SMB demand for AI guidance is accelerating at the exact moment that MSP confidence in delivering it has collapsed, and what that means for the structure of the channel over the next 18 months.

144 Microsoft. "Rise in 'Shadow AI' Tools Raising Security Concerns for UK." Microsoft UK Stories, Oct. 2025.

145 ASUS Business. "The Future of SMB Report 2025: Harnessing the Potential of AI PCs." ASUS, 2025.

146 Woodgate, Scott. "7 Cybersecurity Trends and Tips for Small and Medium Businesses to Stay Protected." Microsoft Security Blog, 31 Oct. 2024.

147 CrowdStrike. "2026 Global Threat Report." CrowdStrike, 2026.

148 Forrester Research. "Predictions 2025: An AI Reality Check Paves the Path for Success." Forrester, 2025.

Part 04.

The Provider Reckoning

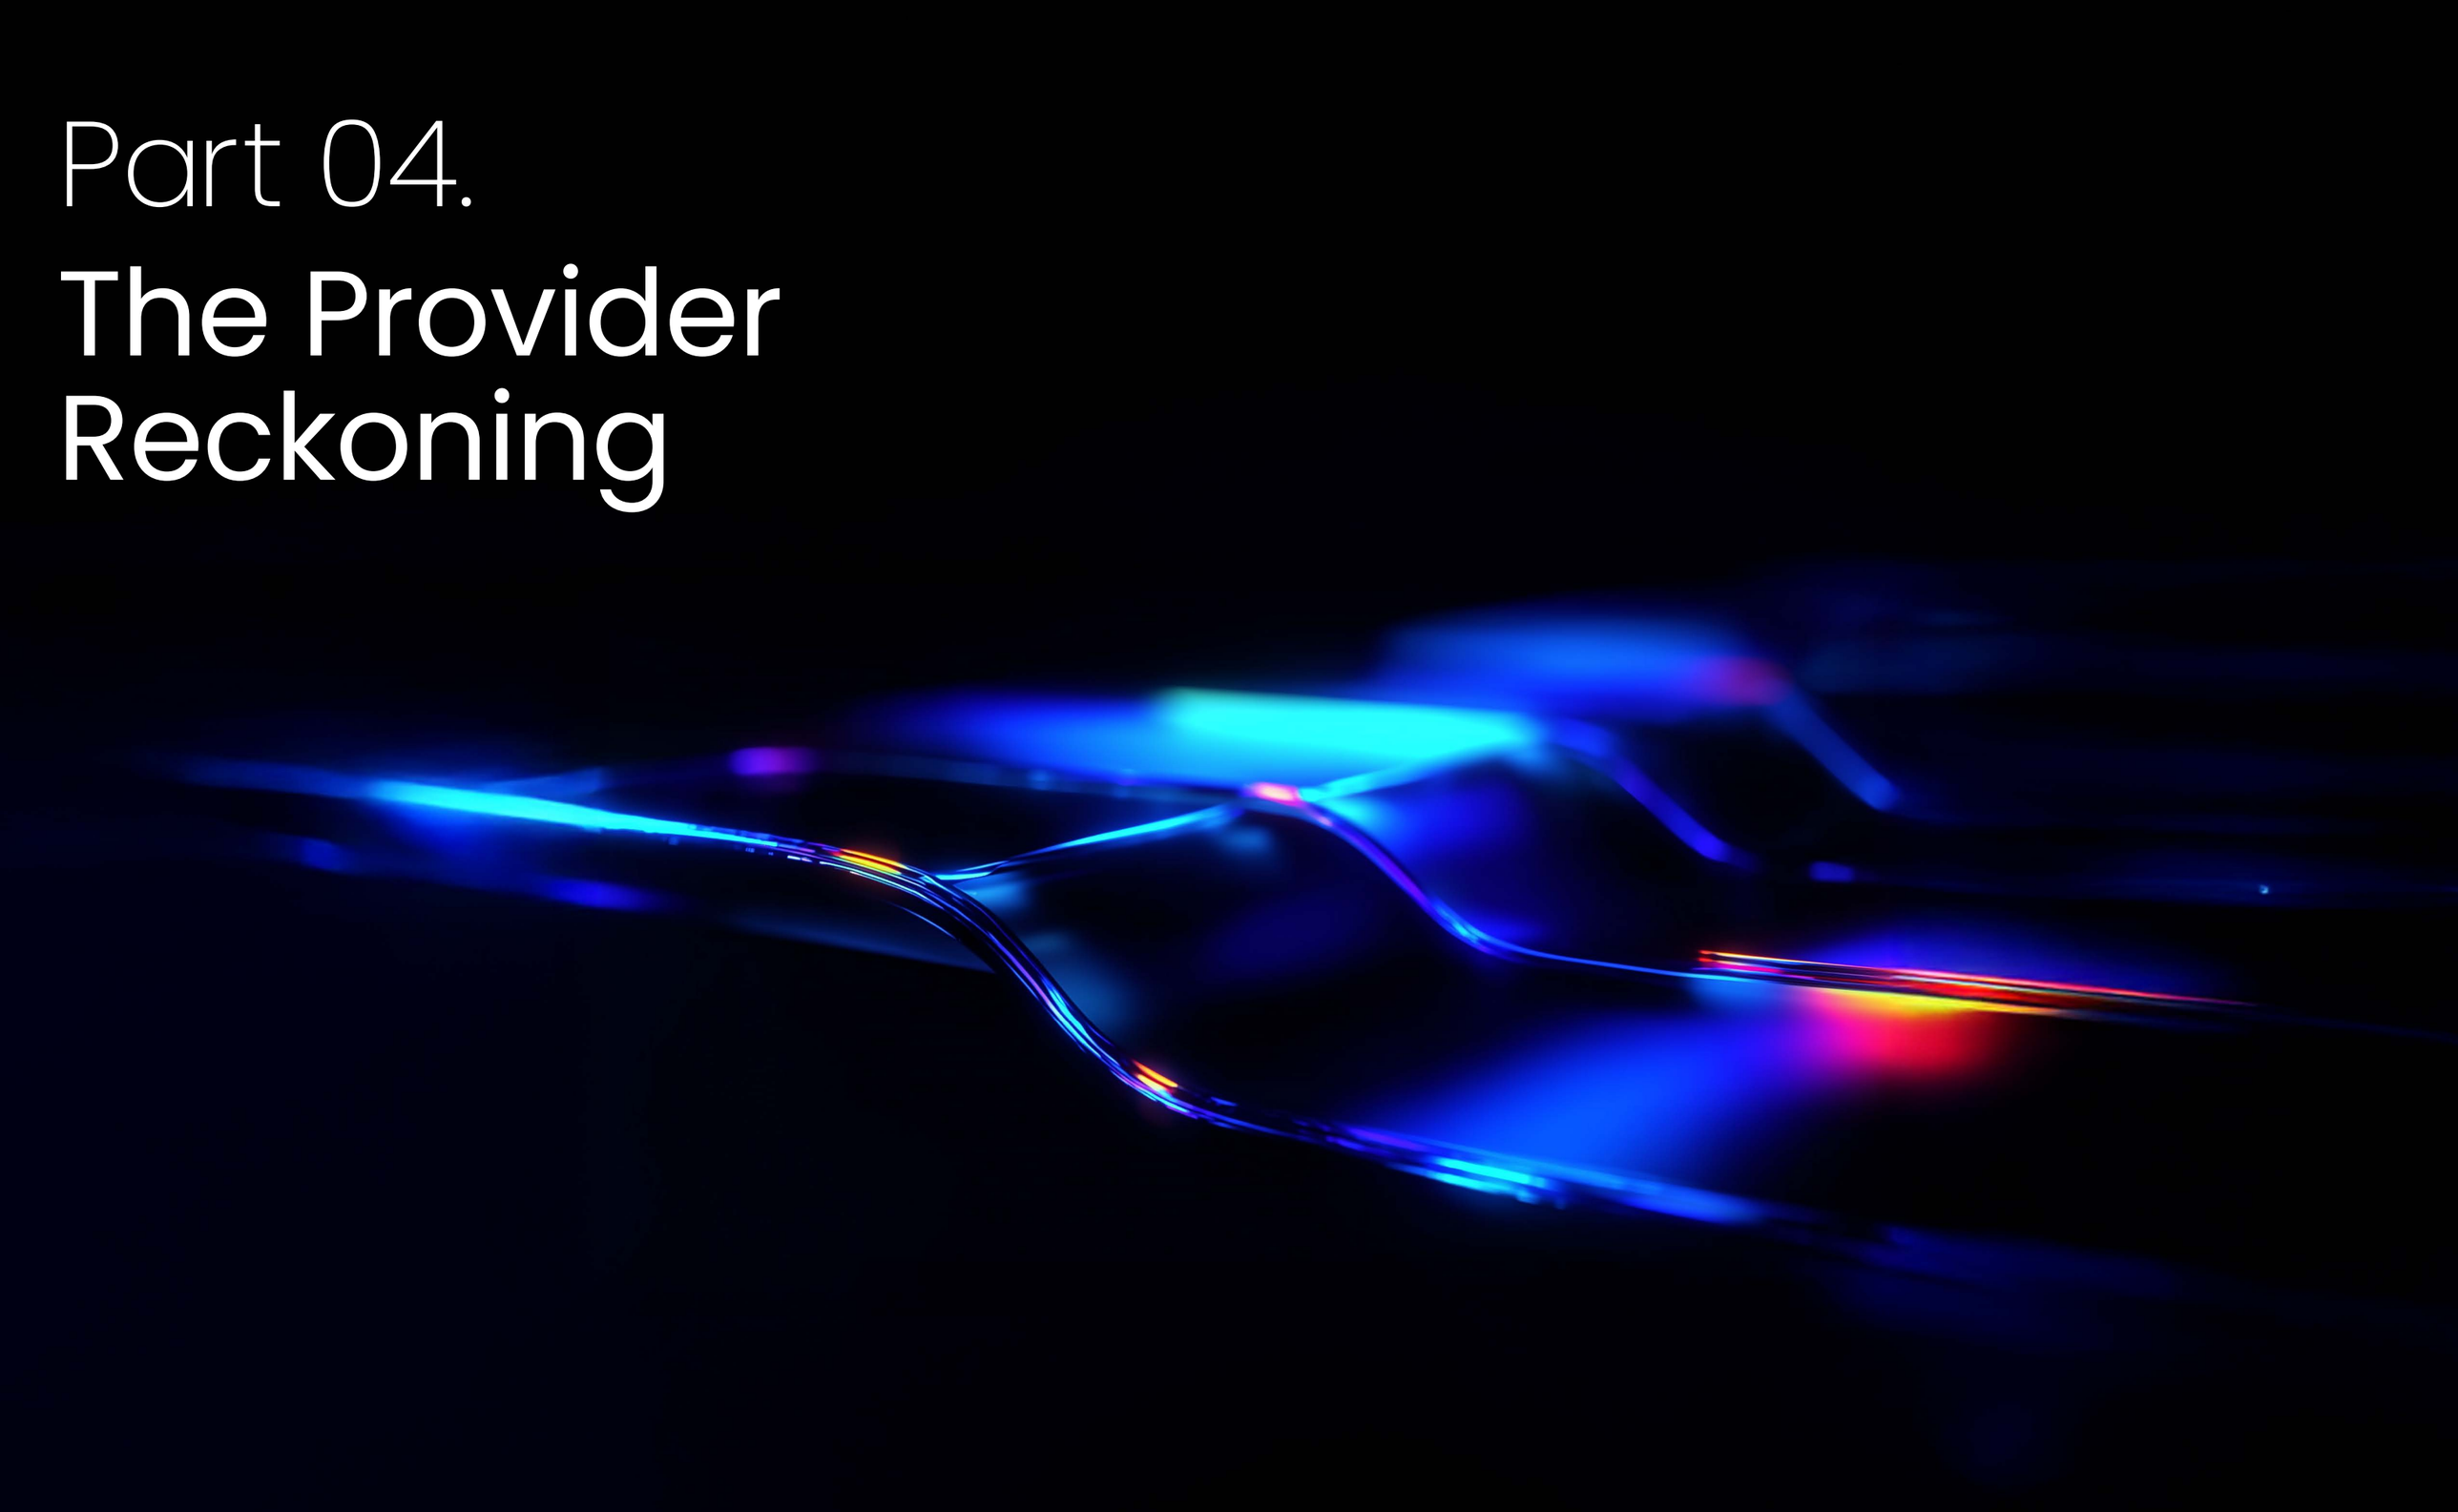

The central finding of the supply-side analysis is one of the more consequential ironies in the current technology market.

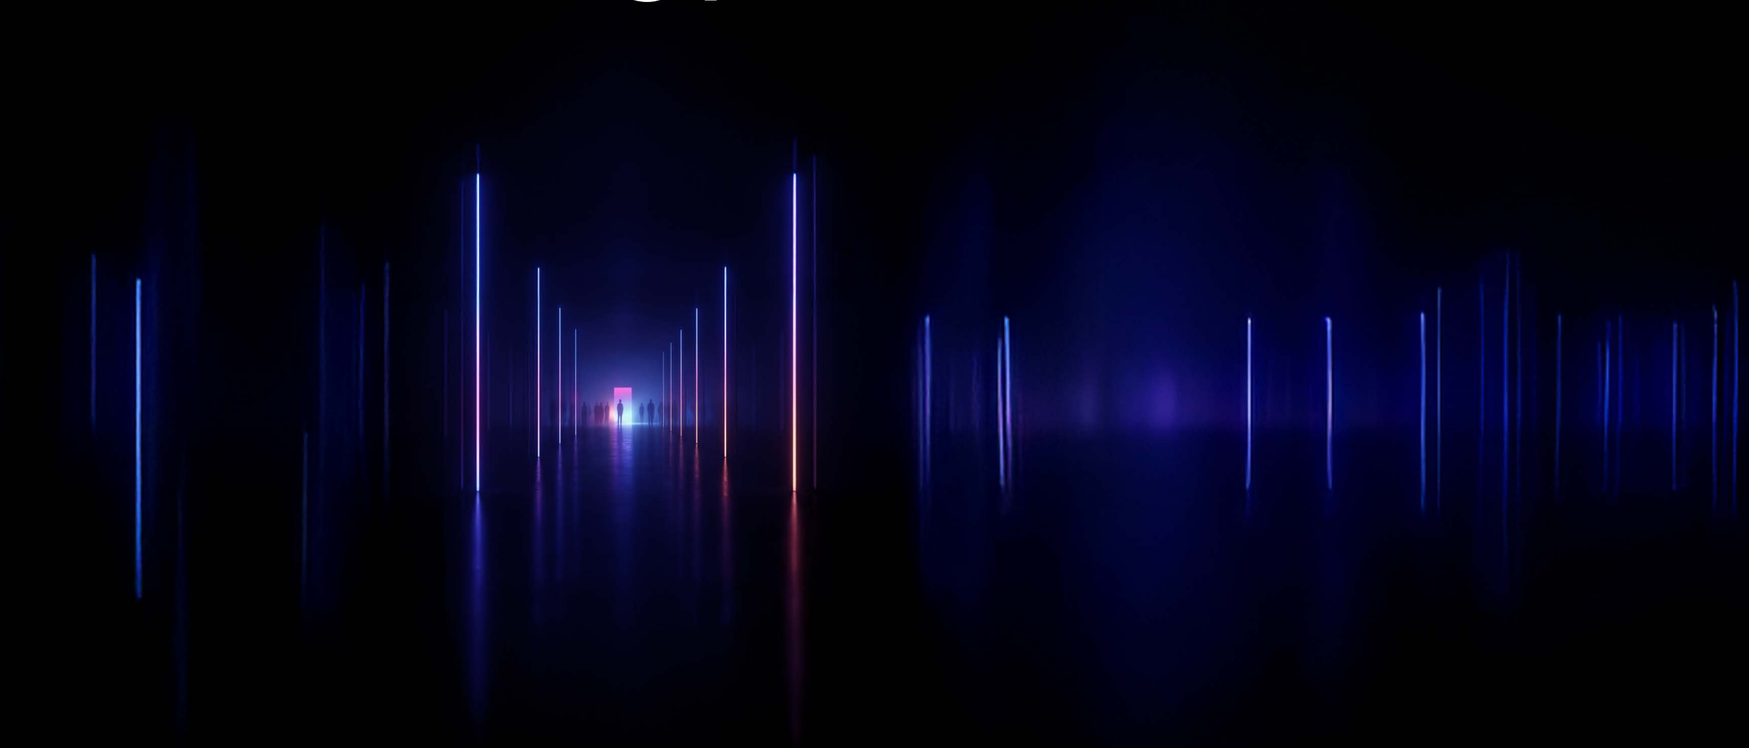

The first three parts of this report have examined the agentic labor economy from the demand side: the macro productivity signal already reshaping the economy, the compounding returns available to businesses that complete the integration journey, the security and data readiness obligations that responsible adoption requires and the structural advantages available to providers who understand all three dimensions as a single, unified challenge. Part Four turns the lens on supply.

The central finding of the supply-side analysis is one of the more consequential ironies in the current technology market: the channel most trusted to guide small and medium businesses through the most complex technology transition of a generation is, by its own admission, losing confidence in its ability to do so at the exact moment that demand is accelerating. That confidence collapse is a structural supply-demand mismatch, and one that is creating a competitive window that will close faster than most of the market currently appreciates.

Part Four examines four dimensions of the provider reckoning now underway. Chapter 10 documents the AI readiness chasm itself: the gap between what SMBs expect from their technology partners and what those partners currently have the capability to deliver. Chapter 11 examines the bifurcation that chasm is accelerating: the splitting of the channel into AI-native challengers and AI-adjacent incumbents, with margin compression as the signal that the restructuring has already begun.

Chapter 12 makes the case for the flywheel strategy that the research most consistently supports: deploying AI internally first, building deployable expertise through operational proof and then externalizing that expertise as a productized service. And Chapter 13 examines the full scope of the managed intelligence opportunity: the $1.3 trillion addressable market, the growth rate differential between AI services and traditional managed services, and the framework that transforms an MSP into a MIP.

Chapter 10.

The AI Readiness Deficit: When Demand Outpaces the Channel

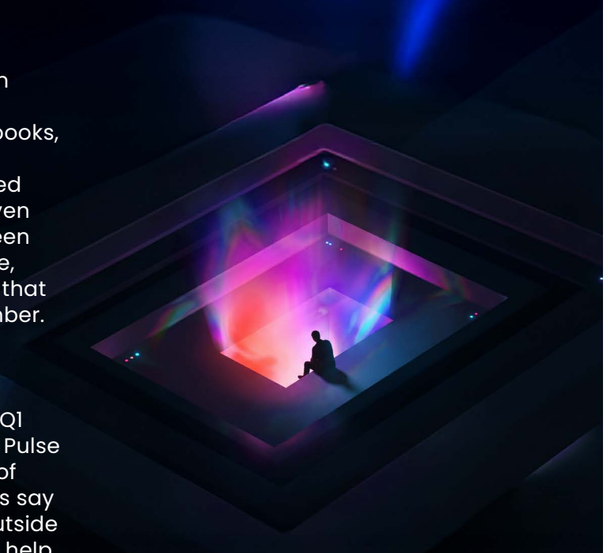

The OpenText Cybersecurity 2025 Global Managed Security Survey of more than 1,000 MSPs contains a finding that should reframe how every technology partner in this market thinks about the competitive environment they are operating in. It is not the headline finding that 92 percent of MSPs reporting that AI is already driving business growth, with 96 percent expecting that trend to continue, though both of those numbers matter.[149] The more consequential finding is the one buried two paragraphs into the press release, the one that is not in any vendor's pitch deck, and the one that describes the structural condition of the channel more accurately than any growth projection can: **in 2024, 90 percent of MSPs felt ready to support AI-related security needs. By 2025, that number had dropped to approximately 50 percent.**[150]

90 percent to 50 percent. In a single year. The confidence collapsed because the gap between what managed service providers were telling clients they could deliver and what they had deployed became impossible to ignore, and the delta was larger than most providers had acknowledged even to themselves. The AI conversation had been happening in the channel for years.

The AI implementation infrastructure; the deployment playbooks, the governance frameworks, the trained practitioners, the proven use cases; had not been built at the same pace, and the survey made that visible in a single number.

The demand signal from the SMB side is unambiguous. Pax8's Q1 2026 SMB Technology Pulse finds that 84 percent of small business leaders say they would trust an outside technology advisor to help their business implement AI, and 94 percent are confident they could find trustworthy outside support to do so.[151] Seventy percent agree that small businesses need outside technology partners to fully benefit from AI.[152] The demand for managed intelligence guidance is already here, and SMBs are actively looking for partners capable of meeting it.

This is the AI readiness chasm. And it is operating on both sides of the provider-client relationship simultaneously, in ways that are beginning to reshape how SMBs select, evaluate and switch technology partners.

149 OpenText Cybersecurity. "OpenText Cybersecurity Finds 92% of Managed Service Providers See AI-Driven Growth, But Readiness Gap Widens." OpenText, 24 Sept. 2025.

150 OpenText Cybersecurity. "OpenText Cybersecurity Finds 92% of Managed Service Providers See AI-Driven Growth, But Readiness Gap Widens." OpenText, 24 Sept. 2025.

151 Pax8. "SMB Technology Pulse Survey: Q1 2026 Topline Summary." Pax8, Q1 2026. Proprietary research. Data on file.

152 Pax8. "SMB Technology Pulse Survey: Q1 2026 Topline Summary." Pax8, Q1 2026. Proprietary research. Data on file.

153 Evans, Katie, et al. "IDC FutureScape: Worldwide Small and Medium-Sized Business 2025 Predictions." IDC, 2024.

154 OpenText Cybersecurity. "OpenText Cybersecurity Finds 92% of Managed Service Providers See AI-Driven Growth, But Readiness Gap Widens." OpenText, 24 Sept. 2025.

155 McCabe, Laurie, and Sanjeev Aggarwal. "The Impact of AI on SMBs: 2025 Trends, Challenges, and Opportunities." SMB Group, June 2025.

THE DEMAND SIDE HAS CROSSED A THRESHOLD

The SMB market did not wait for its technology partners to build AI readiness before developing expectations about it. IDC's research, among the most cited forward-looking datasets in the channel, predicts that 70 percent of SMBs will demand clear, documented AI use cases from vendors and managed service providers before committing new investment in AI or generative AI technologies.[153]

This demand signal indicates that SMBs are actively using AI expertise as a filtering criterion in provider evaluation, as a baseline competency whose absence disqualifies. As the OpenText's survey captures, AI expertise has become the third most important attribute SMBs look for in an MSP, following only threat prevention capability and 24-hour support availability.[154] Those first two attributes have been table stakes for a managed services relationship for a decade. The elevation of AI expertise to that tier; outranking price, responsiveness, tool quality and vertical specialization; signals that the buying criteria for managed services have structurally shifted.

The SMB Group's research, drawing on 650 SMB decision-makers across industries and company sizes, adds texture to what that demand looks like from inside the client organization. Only 30 percent of SMBs describe themselves as very familiar with agentic AI.[155] The overwhelming majority of the market is somewhere on the spectrum between curious and confused. That combination, high demand for AI guidance and low internal familiarity with what agentic AI means in practice, describes a client population that needs a trusted advisor to translate the category for them. The providers who show up with packaged use cases, defined outcomes and clear proof of deployed capability will capture that advisory relationship. The providers who show up with capability decks and partnership logos will not.

THE AI READINESS CHASM: SUPPLY VS. DEMAND TYPE

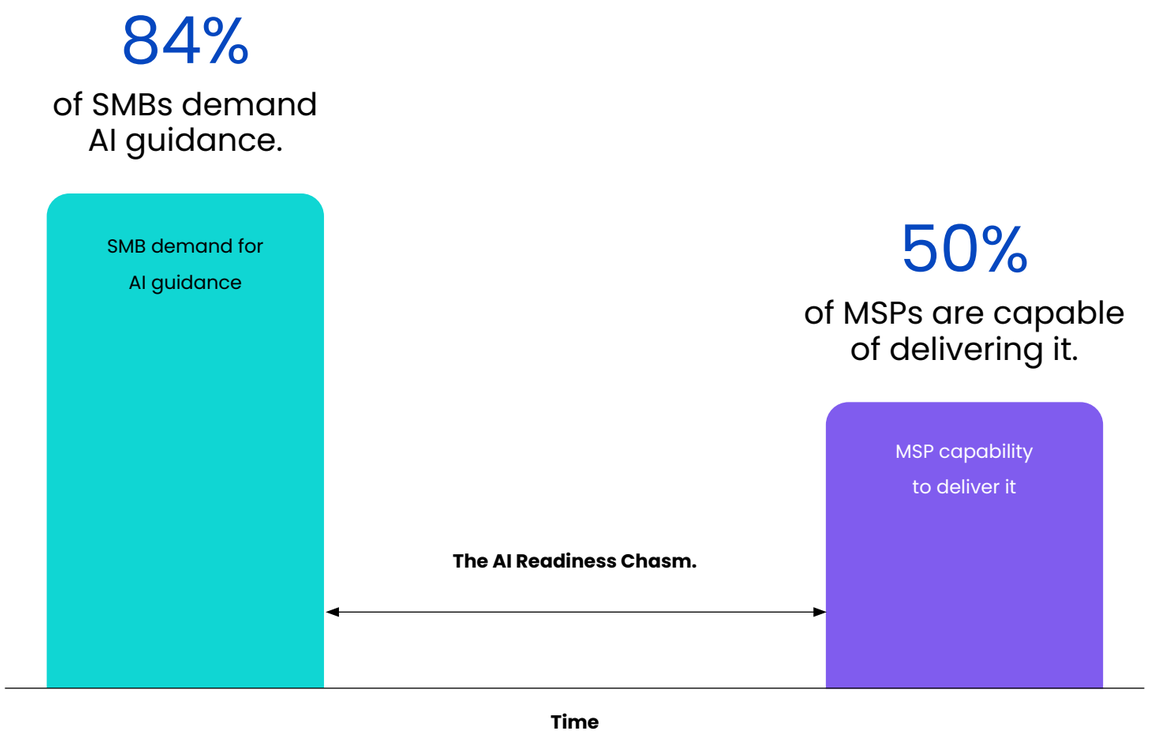

THE SUPPLY SIDE HAS NOT KEPT PACE

The readiness collapse documented by OpenText is not a story about incompetent providers. It is a story about a capable industry that built its expertise on one foundation: reactive security, infrastructure management, compliance monitoring and break-fix resolution. And it's now being asked to perform on a different foundation simultaneously. The providers who built excellent managed security practices did not build those practices in anticipation of needing to advise on agentic AI governance, data architecture for machine learning workflows or the ROI modeling for autonomous agent deployments in SMB environments. No one was building for that five years ago as the need didn't exist.

What makes the chasm structurally significant is that building the missing capability requires time that the market is not providing in abundance. The Omdia research on MSP channel dynamics finds that 61 percent of partners currently struggle to move AI projects beyond proof-of-concept with existing clients; not because the AI tools fail in the proof of concept, but because the partner lacks the workflow redesign expertise, the data infrastructure knowledge and the governance architecture capability to move the project from demonstration to production.[156]

These are organizational capabilities that require deliberate investment, structured learning, and, most importantly, the accumulated experience of having actually deployed AI in production environments where real business processes depend on it.

This is why OpenText's finding that less than half of MSPs have built or deployed AI cybersecurity agents for SMB customers is among the most important supply-side statistics in the current market.[157] It is the clearest available indicator of the distance between the AI conversation most MSPs are having with clients and the AI deployment experience those same MSPs have actually accumulated. The 36 percent experimenting internally with agents but not yet rolling them out to clients are the providers who understand the gap and are trying to close it from the inside.[158] The remainder are operating in a condition that the market is beginning to make untenable.

Pax8's own discovery research provides an unusually direct view of how this readiness collapse is being experienced on the demand side of the relationship. The research interviewed 18 SMB decision-makers across three segments and assessed where each sat on a five-stage AI maturity scale running from "exploring AI but haven't implemented anything" through "AI is the business." Thirteen of the 18 SMBs sat at the same maturity level as the MSPs serving them, somewhere between occasional unstructured AI use and AI embedded in defined operational processes. Among the four tethered SMBs already working with managed service providers, only one described a relationship in which the MSP was actively delivering AI capability into the business. The other three either had no AI conversation with their MSP at all or were the ones pushing the MSP toward adoption rather than the other way around.[159] The implication is structurally consistent with the OpenText supply-side numbers. The MSP channel is not currently leading the AI charge for the SMB economy. It is moving in parallel with it, at roughly the same pace, and the partnership architecture that would convert that parallel motion into directional guidance has not yet been built.

THE INTERNAL ADVANTAGE LEFT ON THE TABLE

The OpenText survey contains within it the most actionable finding for providers who understand the urgency of the chasm: 67 percent of partners already use AI for customer support, 66 percent use it for technical support and ticket triage, and 58 percent use it for threat detection and response.[160] These are operational deployments; AI systems running in production, processing real data, generating real outputs and improving real workflows inside the partners's own service delivery infrastructure.

The providers in those 67, 66, and 58 percent cohorts have something that no certification program, no vendor partnership and no capability deck can replicate: proven, operational AI deployment experience. They know what it takes to integrate an AI tool into a workflow that real staff depend on. They have navigated the data quality problems that emerge when AI systems encounter organizational data for the first time. They have managed the governance questions about which decisions require human review and which can be delegated to the agent. They have measured the actual productivity impact of the deployment and managed the expectation gap when the initial results did not match the initial projections.

Every one of those learnings is directly applicable to a client environment. The AI deployment that reduced ticket triage time inside the MSP's service desk is a template for reducing intake processing time in a professional services firm's billing department. The threat detection agent that monitors the MSP's own infrastructure is the same architecture, governed by the same framework, that can monitor a healthcare SMB's endpoint environment. The providers who have made this connection; who have systematically mapped their internal AI deployments to client workflow equivalents and built the packaging, pricing and delivery mechanisms required to externalize them; are the providers who are closing the readiness gap from the inside out, without waiting for the external training infrastructure to catch up.

THE TRUST VACUUM AND ITS COMPETITIVE CONSEQUENCES

The combination of accelerating client demand and declining provider confidence is not simply a service delivery challenge. It is a client relationship risk with escalating consequences. The SMB that turns to its managed services provider for guidance on AI adoption and receives an uncertain or generic response does not simply defer its AI investment. It begins looking for a different partner — one who can answer the question with evidence rather than aspiration.

ConnectWise's research found that 47 percent of SMBs would switch providers for a stronger cybersecurity solution. [161] **The same switching logic applies with equal force to AI competency.**

An SMB that discovers its current technology partner cannot guide it through the most strategically important technology transition of the decade is not a loyal client exercising patience. It is a client in the early stages of a provider evaluation process, and it is likely to conduct that evaluation quietly, through referrals and conversations at industry events, before the provider is aware the relationship is at risk.

Each of those engagements builds switching costs and deepens the relationship.

The first-mover advantage available to providers who close this gap is about retaining existing ones through a period when the value proposition of the managed services relationship is being re-evaluated against a new set of criteria that most of those clients have not yet articulated out loud. The provider who becomes the trusted AI advisor before the client has formulated a formal AI strategy is the provider who will be embedded in every subsequent AI decision that client makes, the data readiness assessment, the agent deployment, the governance framework, the ongoing performance management. Each of those engagements builds switching costs and deepens the relationship in ways that traditional infrastructure management never achieved.

The providers who wait, who treat AI competency as a capability to build when client demand makes it unavoidable, are deciding whose consequences will not be fully visible until 12 to 18 months from now, when the market has stratified and the clients who drove that stratification have already made their choices.

Chapter 11 examines the shape of that stratification and the margin dynamics that signal how far along it has already progressed.

156 Omdia. "MSP Trends and Predictions 2025 — Executive Summary." Omdia (Canalys), Jan. 2025.

157 OpenText Cybersecurity. "OpenText Cybersecurity Finds 92% of Managed Service Providers See AI-Driven Growth, But Readiness Gap Widens." OpenText, 24 Sept. 2025.

158 OpenText Cybersecurity. "OpenText Cybersecurity Finds 92% of Managed Service Providers See AI-Driven Growth, But Readiness Gap Widens." OpenText, 24 Sept. 2025.

159 Galvan, Moriah. "SMB Discovery." Pax8 UX Research, May 2026. Proprietary research. Data on file.

160 OpenText Cybersecurity. "OpenText Cybersecurity Finds 92% of Managed Service Providers See AI-Driven Growth, But Readiness Gap Widens." OpenText, 24 Sept. 2025.

161 ConnectWise. "SMB Cybersecurity Statistics and Trends." ConnectWise, 2025.

Chapter 11.

The Divide: AI-Native Challengers vs. AI-Adjacent Incumbents

The warning delivered from the mainstage at Kaseya DattoCon 2025 was framed as an observation; a description of a structural condition that was already present in the market and already reshaping the competitive dynamics of the channel. “Margins are under pressure…in many cases, we see that core gross margins have come down from 50 percent to 30 percent. AI-native MSPs have started to enter the market and undercut prices.”[162]

That 20-point compression is the early measurement of a structural economic gap between providers built around human-labor delivery and providers built around automation-first architecture, that will determine which MSPs and MIPs lead the next three years and which find themselves defending margin in a fight the economics have already decided.

Margins are under pressure…in many cases, we see that core gross margins have come down from 50 percent to 30 percent.

TWO TIERS, ONE MARKET

The managed services industry is splitting into two distinct cohorts, and the distinction between them is not primarily one of size, geographic footprint or years in business. It is one of architecture, specifically, whether a provider’s service delivery model is built around AI-augmented operations or built around human labor with AI capabilities added on top.

The AI-native tier is not dominated by startups founded last year. It includes forward-looking established providers who made deliberate, systematic investments in AI-augmented delivery years before the current wave of client demand made those investments urgently necessary. These providers automated their service desks, their monitoring workflows, their patch management processes and their client reporting before their competitors recognized the competitive advantage that automation created. They reduced their cost-to-serve while maintaining or improving service quality. They built the internal AI deployment experience that is now the prerequisite for offering AI services to clients. And they arrived at the current moment of accelerating demand with a cost structure and capability portfolio that allows them to price services in ways that providers still running primarily human-labor delivery models cannot match without margin destruction.

The MSPs who embrace automation and intelligence will lead.

The AI-adjacent tier is larger. It includes much of the channel: providers who understand the direction of travel, who recognize that AI capability is becoming table stakes rather than premium differentiation and who are somewhere on the spectrum between planning to build their AI practice and actively building it. These providers have not made the systematic investments in AI-augmented delivery that would allow them to compete on cost against AI-native entrants. They are competing on relationship, depth of institutional knowledge and breadth of service; advantages that are real and defensible, but that erode predictably as AI-native challengers develop the client intimacy and domain expertise that relationships require.

The other formulation that emerged during DattoCon captures the stakes with the situation demands: “The MSPs who embrace automation and intelligence will lead. Those who don’t will compete on price.”[163] Competing on price against a structurally lower-cost competitor is not a sustainable strategy. It is a path to margin compression that terminates ends in consolidation, acquisition or exit.

162 Succar, Rania. Opening keynote at Kaseya DattoCon 2025, as reported in: Jones, Schyler, et al. “Kaseya CEO Sets Ambitious Vision for MSPs at DattoCon 2025.” MSP Success, Oct. 2025.

163 Succar, Rania. Opening keynote at Kaseya DattoCon 2025, as reported in: Jones, Schyler, et al. “Kaseya CEO Sets Ambitious Vision for MSPs at DattoCon 2025.” MSP Success, Oct. 2025.

THE TOP-EARNER SIGNAL
The clearest signal determining outcomes comes from the ScalePad 2026 MSP Trends Report, which documents a consistent pattern separating top-performing managed services providers from the rest of the market. Top performers are disproportionately likely to offer vCIO and strategic advisory services: 42 percent compared to just 29 percent across all MSPs surveyed and are more likely to invest in customer success practices including technology roadmaps, dedicated account management and structured onboarding. Those investments correlate directly with higher monthly recurring revenue, stronger client satisfaction scores and improved retention rates.[164] The providers earning at the highest levels are doing so in part because they made early investments in advisory infrastructure and AI-augmented delivery, and those investments are now becoming structural advantages that later movers cannot replicate quickly.

The providers earning at the highest levels are earning at those levels in part because they made early investments in the operational infrastructure that AI-native delivery requires; and those investments are now compounding into structural advantages that later movers cannot replicate quickly, regardless of how urgently they recognize the need to try.

The Lansweeper AI in Managed Services research, based on a survey of 195 MSPs across North America and Europe, provides a granular picture of where the broader channel currently sits. 90 percent of MSPs recognize AI as vital to their growth strategy, with nearly two-thirds rating it as "very important."[165] But only 41.5 percent report AI integration levels above 25 percent of their operations, meaning the majority of the channel is in the early stages of adoption at the precise moment that client demand is crossing the threshold from aspirational to requirements-driven.[166] The top barriers telling this story are data quality issues, cited by 93 percent of respondents; legacy system integration challenges, cited by 53 percent; and a shortage of skilled AI professionals, cited by 52 percent.[167] These are not barriers that a software subscription purchase removes. They are organizational capability gaps that require sustained investment and time to close.

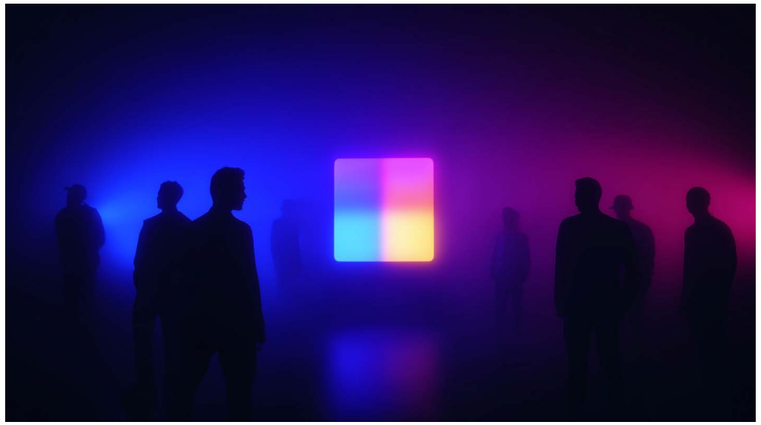

THE MARGIN MATH
The mechanism behind the margin compression described on DattoCon's mainstage is a structural economic advantage. A provider built from the ground up around automation-first delivery, with AI handling tier-one support, monitoring, patch management and client reporting; operates with lower labor intensity, higher scalability per headcount, greater repeatability across engagements and stronger operating leverage than a provider delivering those same services through human effort. As that practice grows, each of those advantages compounds. The AI-native MIP does not need to undercut the incumbent to win the engagement. It simply operates at a fundamentally different cost-per-service-unit; and that difference widens with each new client, each automated workflow and each quarter the incumbent delays building the same architecture.

SMBs are looking for a trusted advisor.

This dynamic plays out first at the commodity end of the service catalog; the foundational services that every managed services provider offers and that most clients treat as undifferentiated. Remote monitoring. Patch management. Help desk support. Backup management. These are the services where AI-native entrants compete most aggressively, where switching costs are lowest and where the margin compression is most acute. The providers who built their practices on these commodity services, who treat them as the foundation of their recurring revenue model, are the most exposed to the bifurcation's first wave.

The providers who survive and lead are the ones moving up the value stack, building practices around the advisory, governance, orchestration and strategic work that AI-native delivery infrastructure enables but does not replace. As Chapter 10 outlined, SMBs are not looking for the lowest-cost provider of remote monitoring. They are looking for a trusted advisor who can translate the implications of the most complex technology transition of a generation into a roadmap they can execute. That work requires the institutional knowledge, client relationships and domain expertise that established providers have built over years of engagement. It is not replicable overnight by an AI-native challenger with a low-cost delivery model but shallow client history.

The path forward for AI-adjacent incumbents is therefore not to compete with AI-native challengers on their own terrain. It is to concede the commodity tier; let it compress, let it consolidate, automate it as aggressively and quickly as possible to reduce internal cost; while simultaneously moving upmarket into the advisory, governance and outcome-based service categories where relationship depth and domain expertise create defensible competitive advantage.

THE 12-TO-18-MONTH WINDOW
ScalePad's 2026 MSP Trends Report confirms the structural divergence already underway: only 39 percent of MSPs currently have an AI roadmap they are actively executing, while the top performers, already differentiated by disproportionate investment in advisory services and customer success infrastructure, are pulling away with each quarter that gap persists.[168] The providers building those structural advantages now are responding to present demand, and compounding their lead in the process. The window to join that cohort rather than be disrupted by it is not indefinitely open.

Market stratification in technology services follows a recognizable pattern. An early adoption phase, where differentiation is available to any provider willing to invest. A growth phase, where the early adopters' advantages compound and late movers begin to recognize the urgency. And a consolidation phase, where the structural separation between tiers becomes self-reinforcing; AI-mature providers attract AI-sophisticated clients, generate the case studies and operational proof that accelerate further client acquisition, invest those returns in deeper capability and widen the gap further with each cycle.

The channel is currently in the transition between the growth phase and the consolidation phase. The providers who close the capability gap in the next 12 to 18 months will arrive at consolidation as leaders. Those who reach that phase still in the AI-adjacent tier will arrive as acquisition targets: providers with valuable client relationships and domain expertise that AI-native acquirers will want, at valuations that reflect the diminishing defensibility of a human-labor delivery model.

Chapter 12 examines the most reliable path from the AI-adjacent tier to the AI-native tier; the internal flywheel strategy that starts not with client-facing AI services but with the MSP's and MIP's own operations, and builds the operational proof, the reduced cost-to-serve and the deployable expertise that transforms a capability gap into a competitive advantage.

164 ScalePad. "MSP Trends Report: See What's Driving Growth in 2026." ScalePad, Jan. 2026.

165 Lansweeper. "AI Adoption in Managed Services." Lansweeper, Mar. 2025.

166 Lansweeper. "AI Adoption in Managed Services." Lansweeper, Mar. 2025.

167 Lansweeper. "AI Adoption in Managed Services." Lansweeper, Mar. 2025.

168 ScalePad. "MSP Trends Report: See What's Driving Growth in 2026." ScalePad, Jan. 2026.

Chapter 12.

Internal First: The Flywheel That Funds the Practice

The most consistent pattern in the research on MSP AI performance is also the most counterintuitive one for providers eager to capitalize on growing client demand: the managed services providers generating the strongest outcomes from AI are not the ones who led with client-facing deployments, but the ones who started with themselves.

This internal-first principle is an operational sequence that the data consistently validates. MSPs and MIPs that deploy AI against their own workflows before deploying it against their clients' workflows gain three compounding advantages simultaneously. They:

- Reduce their own cost-to-serve, building the margin that funds further investment
- Accumulate the real deployment experience — the data quality lessons, the governance discoveries, the workflow redesign insights — that no training program or vendor certification can replicate
- Create the most credible sales asset available in a market where clients are increasingly skeptical of capability claims without evidence.

An MSP or MIP who has reduced their own ticket resolution time by 40 percent through AI automation does not need to promise a client that AI works, because they have already proved that it does.

THE INTERNAL FLYWHEEL

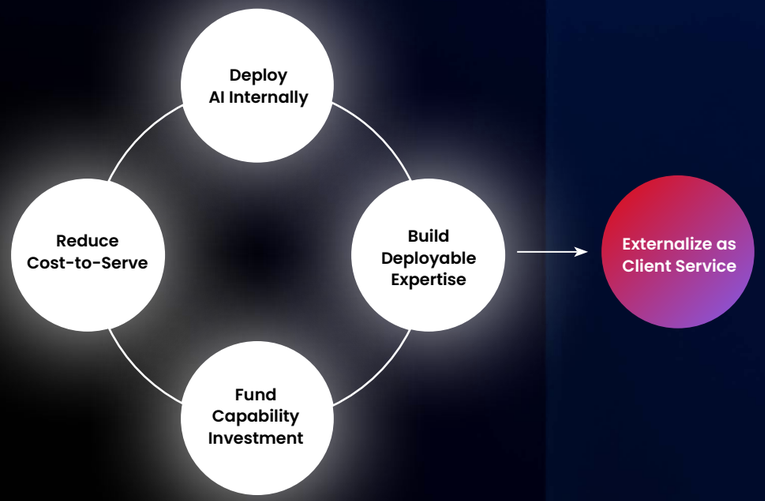

WHAT THE CHANNEL IS ALREADY RUNNING

The OpenText 2025 Global Managed Security Survey, cited earlier in this report, identies the internal deployment reality across the channel: 67 percent of MSPs already use AI for customer support; 66 percent use it for technical support and ticket triage; 58 percent use it for threat detection and response.[169] These are operational deployments running in production, handling real-world workflows, generating real-time performance data and improving service delivery outcomes, quietly, inside the service provider's own operations, largely unconnected to any client-facing AI service offering.

The POPX State of the MSP Industry Survey, conducted among 250 MSP leaders in October 2025, confirms the scale of the operational impact these deployments are generating. 62 percent of MSPs report significant improvements in operational efficiency as a direct result of AI adoption.[170] 86 percent of those using AI-powered chatbots for client support report measurable efficiency improvements.[171] 65 percent report that AI has freed their service agents from repetitive, low-value tasks and redirected that capacity toward higher-value work.[172] These are the kind of operational transformation that, when systematically documented and packaged, becomes the foundation of a differentiated service offering.

The MSP Success Reader Survey, published in December 2025, adds further texture: 93 percent of MSP respondents are already using generative AI tools internally, with operational efficiency improvements spanning executive leadership, help desk and security operations; and 82 percent report AI is delivering the biggest benefits in sales and marketing functions.[173] The channel is not in the early stages of internal AI experimentation. It is, broadly, past experimentation and into operational dependence. The gap is between the providers who have systematically connected that internal dependence to a client-facing service strategy and the providers who have not.

THE FLYWHEEL MECHANISM

The internal-first flywheel operates through three sequential stages, each one funding and enabling the next.

The first stage is cost reduction. MSPs spend up to 80 percent of their total operating costs on labor, according to Omdia's analysis of the MSP cost structure; a figure that reflects the fundamentally human-intensive nature of traditional service delivery.[174] When AI absorbs the highest-volume, lowest-complexity work — tier-one ticket triage, patch management monitoring, routine client communications, documentation generation — it reduces the labor cost per service unit without reducing service quality. MSPs and MIPs deploying AI automation report reducing operational costs by 25 to 40 percent through improved workflow efficiency and reduced manual labor, with help desk resolution times improving by as much as 50 percent where AI triage is deployed at scale.[175] Each point of cost reduction expands the margin available for the deeper capability investments that follow.

169 OpenText Cybersecurity. "OpenText Cybersecurity Finds 92% of Managed Service Providers See AI-Driven Growth, But Readiness Gap Widens." OpenText, 24 Sept. 2025.

170 POPX. "Key Insights from the POPX State of MSP Industry Survey." POPX, Dec. 2025.

171 POPX. "Key Insights from the POPX State of MSP Industry Survey." POPX, Dec. 2025.

172 POPX. "Key Insights from the POPX State of MSP Industry Survey." POPX, Dec. 2025.

173 MSP Success. "AI Is the Next Big Shift: Is Your MSP Ready or at Risk?" MSP Success, Oct. 2025.

174 Omdia. "AI, Ecosystems, and Margin Pressure Have Redefined MSP Strategies in 2025." Omdia (Canalys), Oct. 2025.

175 Acronis. "AI Automation for MSPs: Boost Productivity and Service Quality." Acronis, Jan. 2026.

The second stage is expertise accumulation. This is where the flywheel separates internal-first providers from external-first providers in ways that compound over time and resist replication. Every AI deployment a managed services provider runs against its own operations generates institutional knowledge that cannot be purchased: which data quality problems surface first and how to address them before they corrupt outputs; how to calibrate human-in-the-loop checkpoints so that autonomous execution is safe without becoming a bottleneck; how to design escalation protocols that preserve agent autonomy for routine decisions while routing genuinely complex situations to human judgment; how to measure outcomes in business terms rather than technical metrics so that the value of the deployment remains legible to non-technical stakeholders. A provider who has solved these problems inside their own service desk has solved them in a way that is directly applicable to every client environment with a similar workflow, and the solving takes months of live operational experience, not weeks of vendor training.

The third stage is externalization. The workflows an MSP or MIP has automated internally: ticket triage, monitoring and alerting, documentation, client reporting, patch management, are the same workflows that the provider's SMB clients need to automate. The governance frameworks the provider has built to manage its own AI systems are the same frameworks its clients need to deploy AI agents safely. The outcome metrics the provider has developed to measure AI value in its own operations are the same metrics clients will require before committing to deeper AI investment.

Externalization is not a new product development exercise. It is a packaging and pricing exercise applied to capability that already exists and has already been validated in production.

The Omdia research on the MSP market captures the commercial dimension of what that externalization unlocks. Providers targeting a 30 percent reduction in tier-one and tier-two workloads through AI automation are not only cutting costs; they are also creating the delivery capacity and operational margin that makes AI-augmented client services economically viable to offer at a competitive price point.[176]

The partner who has deployed AI against its own operations can walk a prospective client through specific, measurable, operational outcomes.

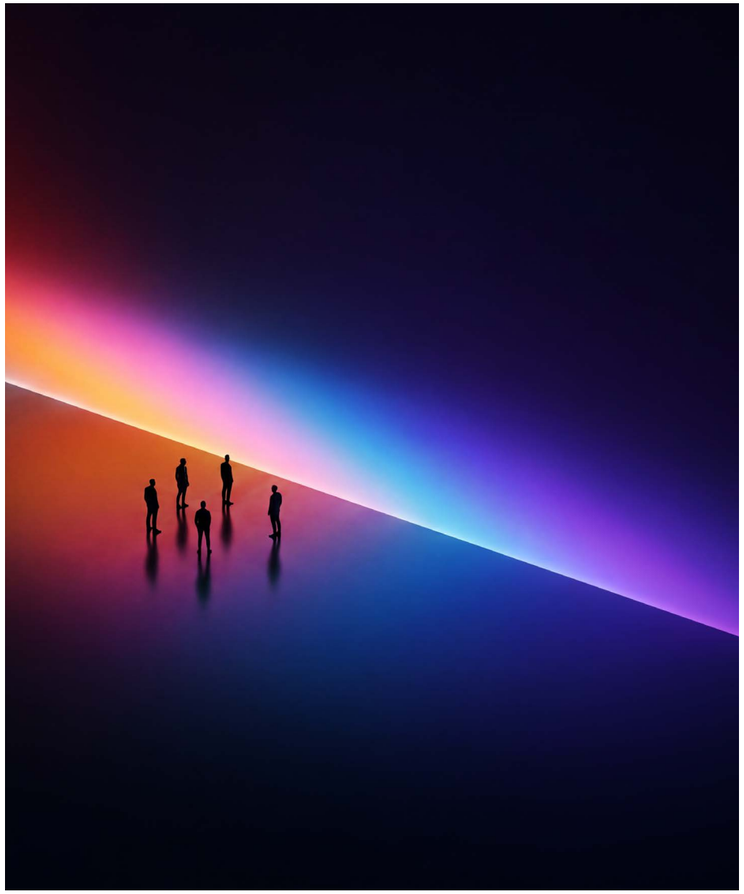

176 Omdia. "AI, Ecosystems, and Margin Pressure Have Redefined MSP Strategies in 2025." Omdia (Canalys), Oct. 2025.

177 Omdia. "AI, Ecosystems, and Margin Pressure Have Redefined MSP Strategies in 2025." Omdia (Canalys), Oct. 2025.

178 POPX. "Key Insights from the POPX State of MSP Industry Survey." POPX, Dec. 2025.

179 Evans, Katie, et al. "IDC FutureScape: Worldwide Small and Medium-Sized Business 2025 Predictions." IDC, 2024.

WHY EXTERNAL-FIRST FAILS

The alternative sequence: leading with client-facing AI services before deploying AI internally, fails consistently. The Omdia October 2025 analysis of MSP AI experimentation found that initial AI deployments frequently stall on the same foundational problems: clean data as a prerequisite for reliable output, workflow redesign as a prerequisite for genuine productivity improvement and governance architecture as a prerequisite for the kind of autonomous operation that generates meaningful time savings.[177] These are not problems that a provider can solve for a client without having solved them first in their own environment. The provider who discovers what "clean data" requires during a client deployment, rather than in their own operations, discovers it at the client's expense in implementation delays, unexpected rework and the credibility damage that comes from promising outcomes that do not arrive on the promised timeline.

The POPX survey finding that 69 percent of MSPs have introduced documented AI strategies or policies to guide responsible use is the governance signature of providers who have run enough internal AI to understand what governance requires.[178] The 31 percent who have not are providers whose AI deployments are either too immature to have generated governance learnings or whose governance gaps remain invisible until a client deployment makes them visible.

69 percent of MSPs have introduced documented AI strategies or policies.

THE CREDIBILITY ASSET

The strategic advantage of the internal-first flywheel that is hardest to quantify, but most significant in practice, is the credibility asset it creates. The managed services market is operating in a period of acute skepticism about AI capability claims. Chapter 10 exposed the confidence collapse at the channel level: 90 percent of MSPs expressing AI readiness in 2024, fewer than 50 percent expressing it in 2025. That skepticism has extended to the client side, where IDC predicts that 70 percent of SMBs will demand proven use cases before committing AI investment.[179]

The partner who has deployed AI against its own operations and can walk a prospective client through specific, measurable, operational outcomes — this is the ticket category we automated, this is the resolution time improvement we measured, this is the governance framework we built to ensure quality, this is what the deployment cost and what it returned — is not engaging in the same sales conversation as the MSP presenting a vendor partnership badge and a capability overview. In a market where evidence is the scarcest and most valuable currency in the AI sales cycle, the internal-first provider has an advantage that is structural rather than circumstantial.

Chapter 13 examines the full scale of the opportunity that the flywheel funds: the transition from MSP to Managed Intelligence Provider, and the $1.3 trillion addressable market that transition positions providers to capture.

Chapter 13.

The MIP Imperative: Managing Intelligence, Not Just Infrastructure

The managed intelligence model is built on orchestrating outcomes.

Every strategic transition in the managed services industry has had a forcing function; a market condition so structurally definitive that providers who fail to respond to it do not merely fall behind but become irrelevant to a category of client need they previously owned. The shift from break-fix to managed services was one such forcing function. The migration to cloud-delivered infrastructure was another. The transition now underway from Managed Service Provider to Managed Intelligence Provider is the third, and the data suggests it will accelerate faster and distribute its consequences more unevenly than either of the transitions that preceded it.

The forcing function this time is not a new delivery model or a new infrastructure paradigm. It is the emergence of agentic AI as the dominant architecture of business technology; a shift in which intelligent agents capable of autonomous reasoning, planning and action replace manual workflows as the primary unit of operational leverage across every business function. The managed services industry built its business model on managing infrastructure. The managed intelligence model is built on orchestrating outcomes. These are not adjacent activities. They require different capabilities, economic models, talent architectures and a fundamentally different relationship with the client.

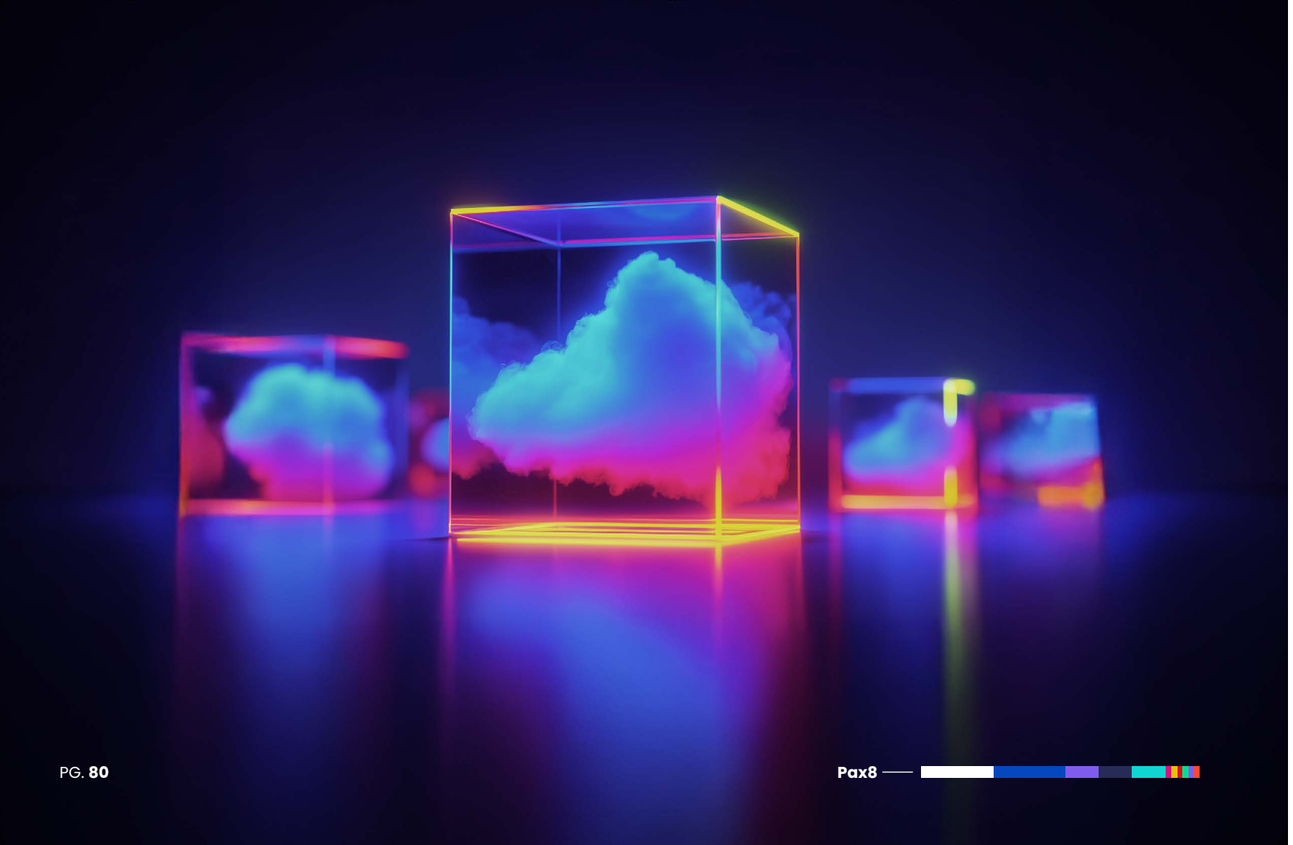

THE SCALE OF THE FORCING FUNCTION

In 2025, IDC published its Worldwide Artificial Intelligence IT Spending Forecast, providing the most authoritative quantification available of the market structural shift now underway. The headline figure is $1.3 trillion in projected AI spending by 2029, growing at a compound annual rate of 31.9 percent: a trajectory that, by 2029, will represent more than 26 percent of all worldwide IT spending.[180] Agentic AI, specifically, drives the most significant component of that growth: an increase in the number and complexity of third-party and custom-built AI agents used by enterprises, a surge in platform spending to enable scalable agent creation and fleet management, and a fundamental reorientation of enterprise IT budgets away from infrastructure maintenance and toward intelligence delivery.

The channel dimension of that forecast is the number that defines the opportunity for managed services providers: through 2029, service providers will account for 80 percent of infrastructure spending as they build out the platforms required to support growing agentic workloads.[181] That forecast is about the primary distribution mechanism through which the largest technology investment wave in the history of the global economy will flow. The channel is not adjacent to this transition. It is structurally central to it.

Analysts at Canalys similarly captured the growth rate differential that translates this structural position into a strategic imperative in terms every managed services provider should internalize: AI services are growing at 59 percent, while managed services are growing at 13 percent.[182] **A provider growing its AI services practice is growing at four and a half times the rate of a provider growing its traditional managed services practice.** A provider with no AI services practice is growing at one-fifth the rate of the market opportunity it is positioned to serve. The compounding effect of that differential, sustained over 36 months, is the divergence Chapter 11 explored and the consolidation it predicts.

$1.3 trillion

in projected AI spending by 2029

180 IDC. "Agentic AI to Dominate IT Budget Expansion Over Next Five Years, Exceeding 26% of Worldwide IT Spending, and $1.3 Trillion in 2029." IDC, 26 Aug. 2025.

181 IDC. "Agentic AI to Dominate IT Budget Expansion Over Next Five Years, Exceeding 26% of Worldwide IT Spending, and $1.3 Trillion in 2029." IDC, 26 Aug. 2025.

182 McBain, Jay. "Partners Less Confident in Channel Revenue Growth." Channel Futures, Mar. 2025. Citing Canalys channel growth rate analysis.

WHAT IS THE MIP MODEL?
The term "Managed Intelligence Provider" was introduced by Pax8 in its June 2025 research report, *The Agentic Inflection Point*, and operationalized through its October 2025 *Managed Intelligence Provider Playbook*: the most substantive channel-facing framework yet published for understanding what the transition from MSP to MIP actually requires in practice.[183] The framework deserves examination not as a structural definition of what the market is demanding but rather what providers must be able to deliver.

The core distinction the MIP framework draws is between managing infrastructure and orchestrating intelligence. Managed service providers, in the traditional model, manage uptime, availability and the operational continuity of technology systems their clients depend on. That model is not disappearing, but rather, collapsing into the minimum viable offering that every provider in the market will deliver and no provider will differentiate on. What clients are now asking for, and what the Pax8 research found through its partner survey data, is something categorically different: strategic guidance on AI adoption, help discovering and managing AI agents, support integrating digital labor into existing workflows and outcome-based service delivery measured against business results rather than system availability metrics.[184]

The shift in partner self-perception captured in the Pax8 survey data is a leading indicator of how rapidly this transition is already reshaping provider identity. Among Pax8 partners surveyed, 66 percent believe they will be seen by their clients as strategic business advisors within two years; a more than seven-fold increase from their current self-identification as IT support providers.[185] That is a recognition, already embraced among forward-looking providers, that the client relationship is being redefined and that the providers who succeed in the next era will be those who complete the transition from reactive vendor to proactive orchestrator before their clients begin evaluating alternatives.

The term "Managed Intelligence Provider" was introduced by Pax8 in June 2025.

183 Pax8. "Pax8 Releases 'The Managed Intelligence Provider Playbook' to Empower MSPs in the Agentic Economy." Pax8, 6 Oct. 2025. Complementing: Pax8. "Pax8 Introduces the Era of Managed Intelligence in its 2025 Research Report: The Agentic Inflection Point." Pax8, 11 Jun. 2025.

184 Pax8. "Pax8 Releases 'The Managed Intelligence Provider Playbook' to Empower MSPs in the Agentic Economy." Pax8, 6 Oct. 2025. Complementing: Pax8. "Pax8 Introduces the Era of Managed Intelligence in its 2025 Research Report: The Agentic Inflection Point." Pax8, 11 Jun. 2025.

185 Pax8. "Pax8 Releases 'The Managed Intelligence Provider Playbook' to Empower MSPs in the Agentic Economy." Pax8, 6 Oct. 2025. Complementing: Pax8. "Pax8 Introduces the Era of Managed Intelligence in its 2025 Research Report: The Agentic Inflection Point." Pax8, 11 Jun. 2025.

186 Pax8. "Pax8 Releases 'The Managed Intelligence Provider Playbook' to Empower MSPs in the Agentic Economy." Pax8, 6 Oct. 2025. Complementing: Pax8. "Pax8 Introduces the Era of Managed Intelligence in its 2025 Research Report: The Agentic Inflection Point." Pax8, 11 Jun. 2025.

187 MSP Success. "AI Is the Next Big Shift: Is Your MSP Ready or at Risk?" MSP Success, Oct. 2025.

188 Pax8. "Pax8 Releases 'The Managed Intelligence Provider Playbook' to Empower MSPs in the Agentic Economy." Pax8, 6 Oct. 2025. Complementing: Pax8. "Pax8 Introduces the Era of Managed Intelligence in its 2025 Research Report: The Agentic Inflection Point." Pax8, 11 Jun. 2025.

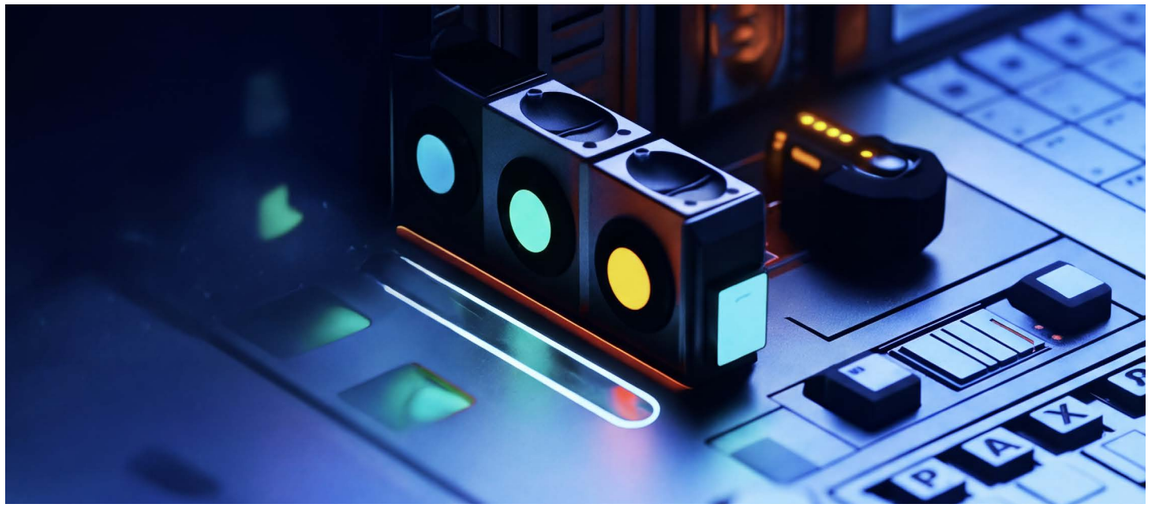

THE SIX PLAYS OF THE MIP TRANSITION
The MIP Playbook defines six capabilities: Discover, Sell, Buy, Build, Implement and Manage; that describe the functional architecture of a managed intelligence practice.[186] Each play represents a category of client engagement that a traditional managed services provider does not currently own and that an MIP must be able to deliver credibly to capture and retain the advisory relationship.

The Discover Play begins not with products but with process; helping clients identify the manual, fragmented and repetitive work that drains productivity and represents the highest-leverage opportunity for agentic intervention. The most common client request to MSPs, according to the MSP Success survey, is help with AI automation and workflow optimization, cited by 64 percent of respondents.[187] The providers who can walk a client through a structured discovery of their process debt and translate that audit into a prioritized deployment roadmap are the providers who own the entry point to every subsequent AI investment that client makes.

The Sell Play involves externalizing internally proven capabilities as packaged client-facing services: the flywheel mechanism outlined in Chapter 12. The Buy Play involves curating and deploying pre-vetted agents and AI-enhanced software from trusted sources; becoming the trusted filter through which an SMB client accesses an agent marketplace that is expanding too rapidly for most organizations to evaluate independently. The Build Play involves developing custom agentic workflows where off-the-shelf solutions do not address client-specific requirements. The Implement Play involves organizational change management and consistent upskilling and reskilling. And the Manage Play involves the ongoing optimization, compliance management and lifecycle oversight that ensures agents remain secure, effective and aligned with evolving business requirements over time.

Together, these six Plays define a service architecture that is outcome-oriented rather than infrastructure-oriented, advisory rather than reactive and economically structured around the value of intelligence delivered rather than the cost of labor consumed. The pricing models that accompany this architecture: per agent, per action, per workflow, per outcome, replace the per-seat and per-device models that have defined managed services economics for two decades with models that align provider revenue with client value.[188]

Not every partner enters this architecture at the same point, and the MIP model is designed to accommodate that. The Broker track offers a lower-barrier entry: deploying pre-built agent configurations, activating AI capabilities already embedded in existing vendor stacks and delivering governed outcomes without requiring custom development capability. The Builder track is for partners ready to design and deploy custom agent workflows, build proprietary vertical playbooks and operate at the full scope of the managed intelligence relationship. Both tracks lead to the same destination. The Broker entry point is not a lesser version of the MIP; it is the starting position from which most partners will build, and it generates real client value and recurring revenue from day one. The architecture meets partners where they are and grows with them as their capability compounds.

THE MIP PRACTICE ARCHITECTURE.

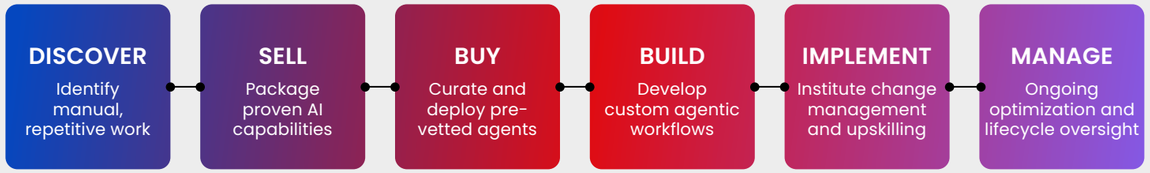

The partner who owns token consumption owns a billing relationship that is as recurring, as mission-critical and as difficult to displace as any infrastructure contract in the traditional model.

THE NEW ECONOMIC MODEL

The economic model underlying the MIP transition is more favorable to providers who complete it than the traditional managed services model it replaces. The traditional model prices labor and availability. The intelligence model prices outcomes and orchestration. A provider billing for the outcome of a workflow automation; accounts receivable processing time reduced by 40 percent, ticket resolution time reduced by 50 percent, client onboarding time cut by 25 percent; is billing for demonstrated business value at a price point that reflects that value rather than the cost of the human labor it replaced. The margin structure of intelligence services is structurally superior to the margin structure of infrastructure management, and it produces the client stickiness that infrastructure management increasingly cannot sustain as AI-native challengers commoditize the delivery tier.

There is a new line item on the SMB balance sheet that most technology partners have not yet built into their client conversations: token consumption. Every business running an agentic workforce pays for that workforce in tokens: the unit of AI computation that governs how much a model processes, reasons and outputs on a given task. As agentic systems scale, token consumption does not scale linearly. Agents running overnight, executing loops, coordinating across domains simultaneously and maintaining persistent memory generate token volumes that look nothing like the predictable, session-based usage patterns of conversational AI tools. Seventy-eight percent of IT leaders already report unexpected charges from consumption-based AI pricing models, and 90 percent of CIOs cite cost forecasting as their top challenge in AI deployment.[189] Those figures reflect early, relatively modest agentic deployments. As SMB clients scale their digital workforces, the gap between budgeted and actual token spend will widen, and the technology partner who has not built token governance into their service model will be explaining surprises rather than preventing them.

The managed intelligence provider who governs token consumption: monitoring agent activity, optimizing prompt architecture, managing model selection across task types and forecasting consumption against budget, is providing a service that has no equivalent in the traditional managed services catalog. Call it token-roll: the ongoing management of the AI workforce's operating cost, sitting alongside payroll as the second workforce budget line every agentic SMB will carry. The partner who owns that function owns a billing relationship that is as recurring, as mission-critical and as difficult to displace as any infrastructure contract in the traditional model, and that compounds in value as the client's agentic workforce grows.

Gartner's projection that global AI services spending will reach $644 billion in 2025 alone; cited by Pax8 as the backdrop against which it released the MIP Playbook; frames the scale of the addressable market that outcome-based intelligence services can capture.[190] The providers positioned to capture their share of that market are not those waiting for the transition to become unavoidable. They are those who have already deployed AI against their own operations, built the expertise and governance frameworks required to extend those deployments to client environments and begun packaging their internal proof into the kind of proven, outcome-measured service offering that the IDC research predicts 70 percent of SMBs will require before committing new AI investment.[191]

189 Flexprice. "Why AI Companies Have Adopted Usage-Based Pricing in 2026." Flexprice, Feb. 2026.

190 Pax8. "Pax8 Releases 'The Managed Intelligence Provider Playbook' to Empower MSPs in the Agentic Economy." Pax8, 6 Oct. 2025.

191 Evans, Katie, et al. "IDC FutureScape: Worldwide Small and Medium-Sized Business 2025 Predictions." IDC, 2024.

192 Pax8. "Pax8 Releases 'The Managed Intelligence Provider Playbook' to Empower MSPs in the Agentic Economy." Pax8, 6 Oct. 2025. Complementing: Pax8. "Pax8 Introduces the Era of Managed Intelligence in its 2025 Research Report: The Agentic Inflection Point." Pax8, 11 Jun. 2025.

193 Pax8. "SMB Technology Pulse Survey: Q1 2026 Topline Summary." Pax8, Q1 2026. Proprietary research. Data on file.

194 Galvan, Moriah. "SMB Discovery." Pax8 UX Research, May 2026. Proprietary research. Data on file.

THE WINDOW

The managed services industry has navigated major transitions before. The break-fix to managed services transition rewarded the providers who moved early with recurring revenue models and long-term client relationships that late movers had to acquire at premium multiples. The on premises to cloud transition rewarded the providers who built cloud competency before their clients demanded it with advisory relationships that shaped every subsequent infrastructure decision those clients made.

The MSP-to-MIP transition will follow the same pattern, compressed into a shorter window by the pace at which agentic AI is penetrating the SMB market. The CEO of Pax8 provided his own assessment: the rate of advancement makes the next 24 to 36 months critical, and MSPs who evolve quickly into the MIP model will be positioned for long-term success while those who delay will find the model shift permanent.[192] This is a structural observation about a market in which the top-earning providers are already differentiating on advisory and outcome delivery, the confidence gap between what clients expect and what providers can deliver is widening by the quarter and the AI-native challengers from Chapter 11 are building cost structures that make the commodity tier of managed services economically untenable for providers who remain there.

The urgency driving that demand is not coming uniformly from SMB leadership. Pax8's Q1 2026 SMB Technology Pulse finds that operational leaders, those closest to day-to-day workflows, are significantly more likely than owners and founders to believe AI will be required for competitiveness: 70 percent versus 56 percent. Seventy-three percent of functional leaders say their business must move on AI within six months.[193] The people inside SMB organizations who understand what AI can absorb are already driving the conversation. The MIP who reaches them, not just the owner, is the one who owns the engagement.

Pax8's discovery research adds an uncomfortable detail to the timeline this chapter has been describing. Among the SMBs interviewed who were currently tethered to a managed service provider, the dominant pattern was that AI adoption was happening around the MSP relationship rather than through it. The tools the SMB was experimenting with had been chosen by the SMB. The agentic use cases were being scoped by the SMB. The strategic conversations about where AI fit into the business were happening inside the SMB and, in some cases, with the MSP only after the SMB had already initiated them.[194] For a managed service provider planning a deliberate MIP transition, that pattern is the strategic warning embedded in the broader window argument. The transition is not simply a matter of building MIP capability before client demand arrives. It is a matter of building that capability before the client builds an AI strategy that does not include the provider. Each quarter that an MSP defers the MIP transition is a quarter in which existing clients are deepening AI relationships the MSP is not party to, with vendors, internal champions and adjacent advisors who will own the next conversation by virtue of having owned the last one.

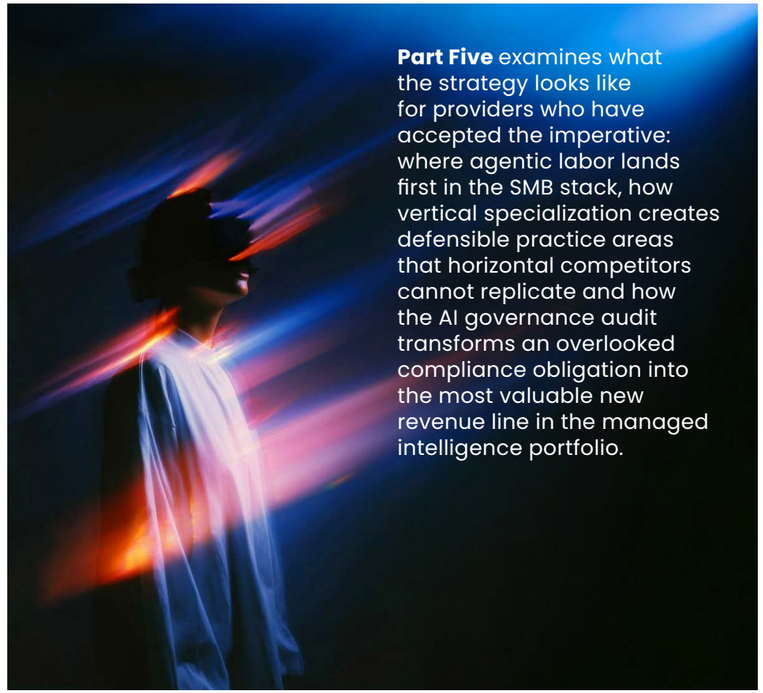

Part Five examines what the strategy looks like for providers who have accepted the imperative: where agentic labor lands first in the SMB stack, how vertical specialization creates defensible practice areas that horizontal competitors cannot replicate and how the AI governance audit transforms an overlooked compliance obligation into the most valuable new revenue line in the managed intelligence portfolio.

Part 05.
The Strategy

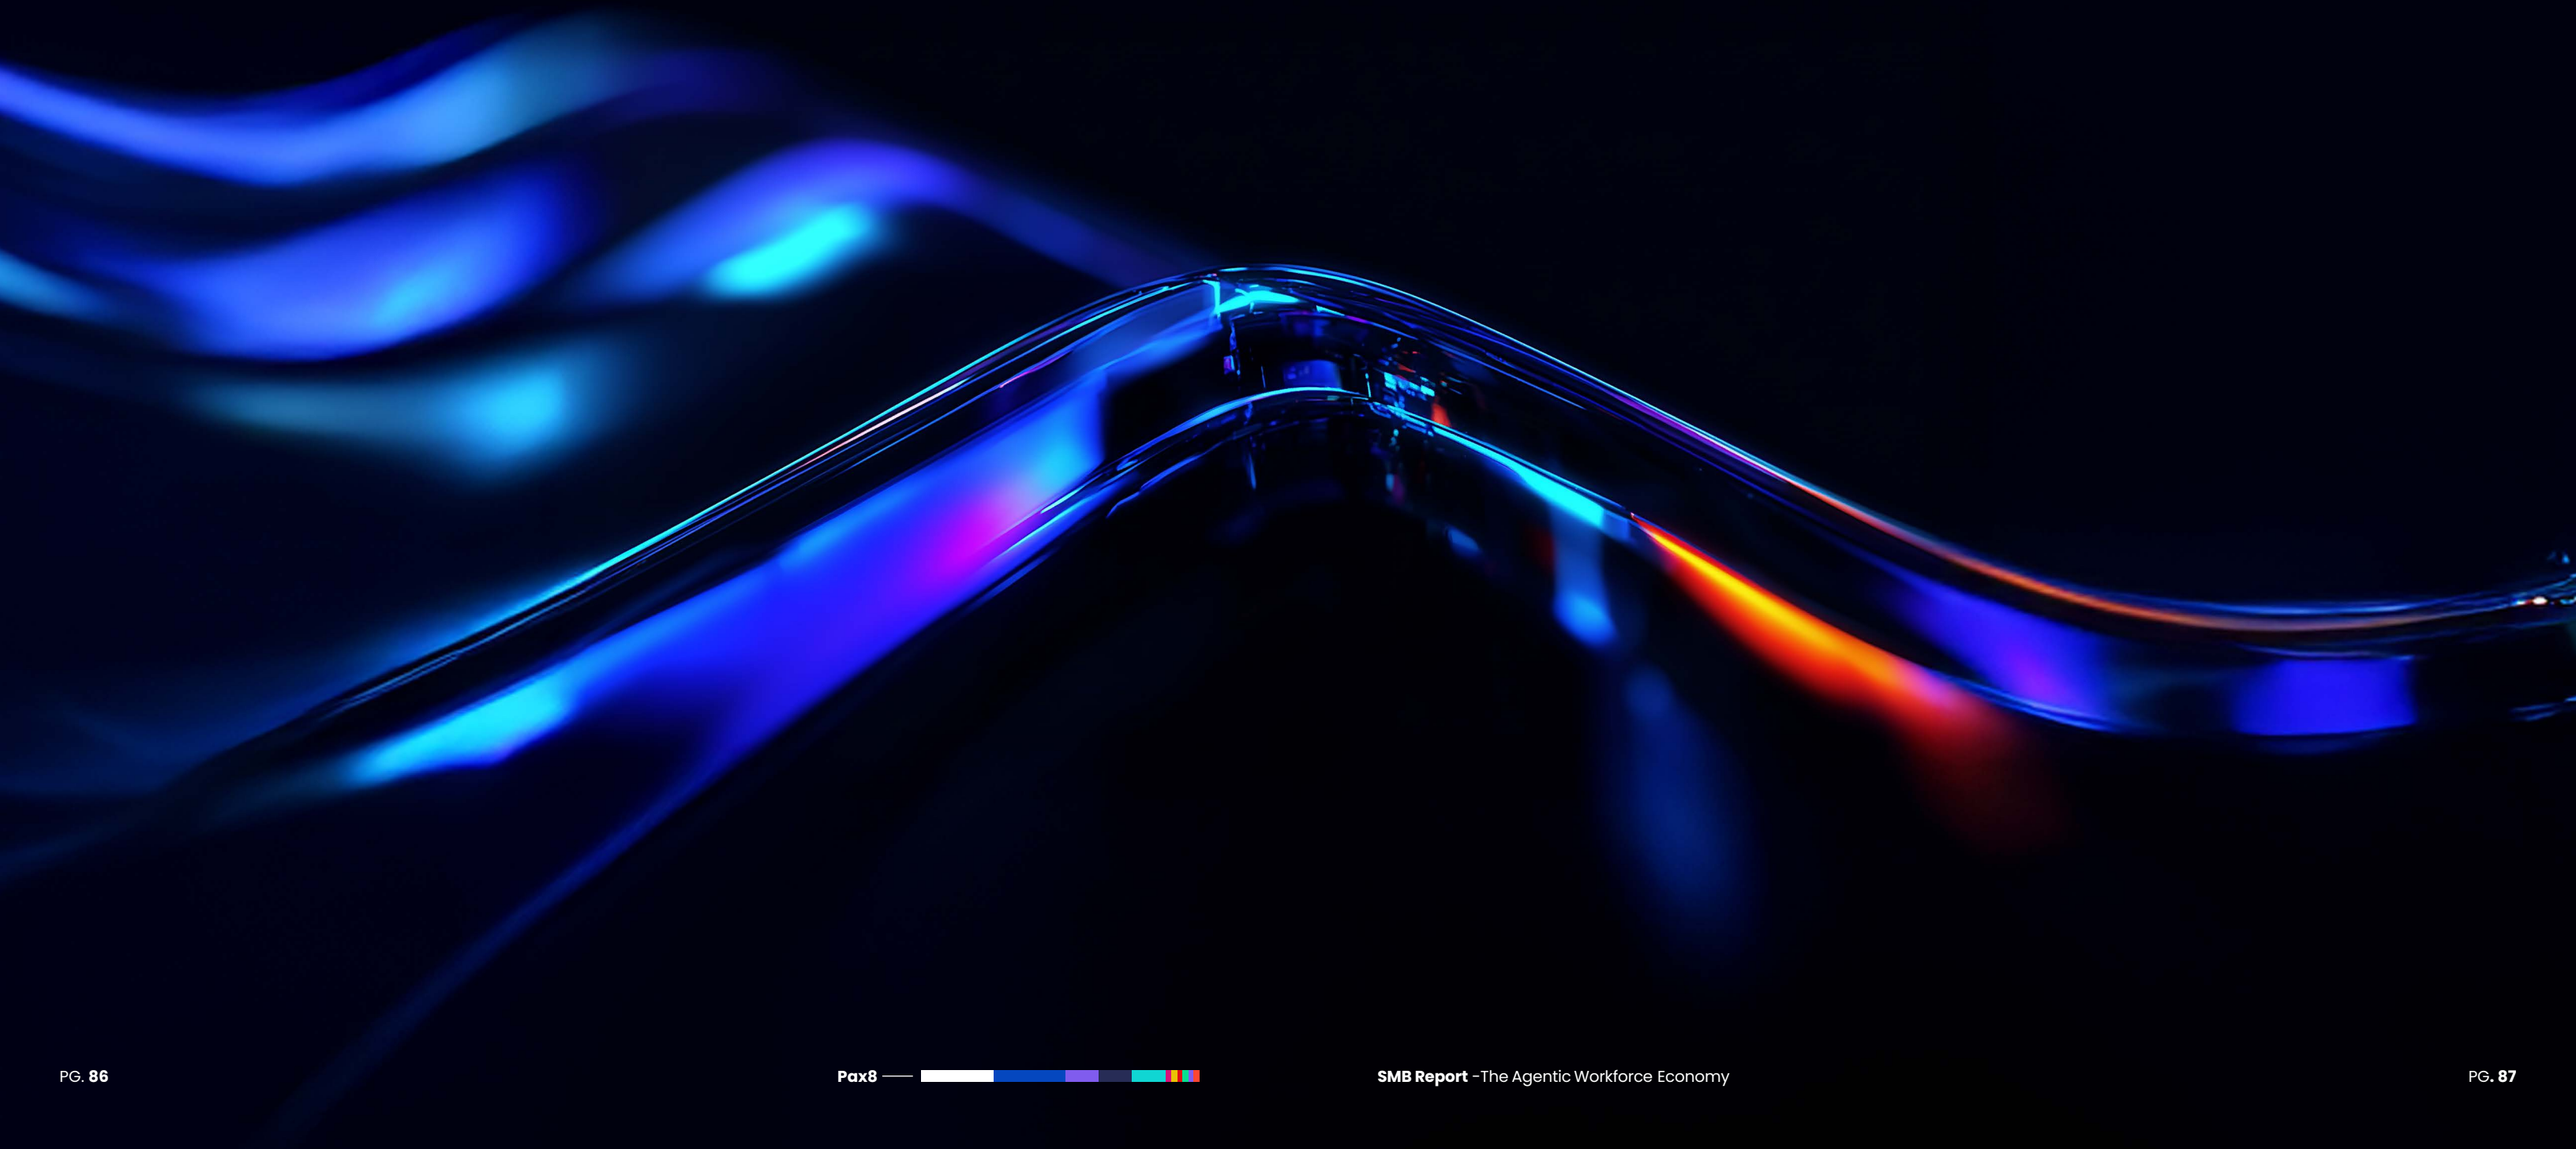

Most SMB clients are already running AI tools their technology partners do not know about.

The four parts of this report have built a case from the ground up. Part One established the economic foundation: the productivity signal already reshaping output across the economy, the nonlinear return curve that rewards businesses that complete the integration journey, the time dividend that most SMBs are recapturing but not yet reinvesting and the structural advantage available to businesses that build on AI-native architectures from the start. Part Two examined the risk dimension: the cybersecurity exposure that grows in direct proportion to agentic labor adoption, the data readiness deficit that undermines AI performance while creating new compliance vulnerabilities and the governance gap that represents both the greatest threat to agentic AI deployments and the most significant untapped opportunity for technology partners willing to fill it. Parts Three and Four turned the lens on the channel: documenting the AI readiness chasm that has opened between what SMBs expect from their technology partners and what those partners currently have the capability to deliver, the bifurcation that chasm is accelerating between AI-native challengers and AI-adjacent incumbents, the flywheel that providers who deploy AI internally first are using to fund their transition and the full scope of the managed intelligence opportunity that transition positions them to capture.

Part Five examines the question that all four of the preceding parts converge on: where does the strategy begin?

Not the strategy in the abstract sense of market positioning and service portfolio design; Parts Three and Four addressed those dimensions in detail. The strategy in the operational sense: where does the first conversation happen, what does the first engagement look like and how does the provider sequence the work from initial deployment to compounding practice. The following chapters address three specific strategic execution questions:

Chapter 14 examines where agents land first in the SMB stack; the finding that digital labor adoption in the small business economy follows a consistent operations-first pattern driven by where digital infrastructure already exists, pain is already established and baselines are already measurable. Understanding this pattern is the difference between positioning agentic labor in a way that generates immediate traction and positioning it in a way that stalls before it starts.

Chapter 15 makes the case for vertical playbooks over horizontal platforms; the finding that generic AI capabilities lose to industry-specific agent workflows with predefined metrics, and guardrails and that the providers who package their domain expertise into repeatable, fixed-scope deployment frameworks break through pilot purgatory while horizontal competitors stall.

Chapter 16 examines the AI governance audit, the most overlooked and most immediately actionable revenue opportunity in the managed intelligence portfolio. Most SMB clients are already running AI tools their technology partners do not know about. The governance audit: discovering ungoverned AI, assessing risk, integrating tools into managed workflows and establishing policy frameworks, is higher-margin, stickier work than initial deployment and it starts with clients the provider already has.

Chapter 14.

The Digital Labor Stack: Where Agents Land First

The most common mistake technology partners make when positioning agentic labor to SMB clients is starting with the wrong question. They ask: what is AI capable of doing? The better question and the one that predicts where adoption occurs is: where does a small business already have agentic infrastructure and a measurable baseline against which to judge improvement?

Where does a small business already have agentic infrastructure and a measurable baseline against which to judge improvement?

Those two questions produce very different answers. The first generates a list of impressive AI capabilities, many of which require data architecture, workflow redesign and organizational change management that most small businesses are not ready to undertake. The second generates a much shorter, more prosaic list: payment collection, billing and invoicing, client information management, scheduling, and where the foundations required for agent deployment already exist. Understanding this distinction is the difference between positioning digital labor in a way that generates immediate traction and positioning it in a way that stalls in proof-of-concept indefinitely.

WHAT THE REAL SMB DIGITAL BEHAVIOR SHOWS

The vcita and inTandem 2025 SMB survey, drawing on behavioral data from hundreds of thousands of small business users, provides one of the most grounded pictures available of where digital tool adoption in the small business economy has concentrated. The top tools SMBs report using are, in order: payment collection tools (used by 41 percent of respondents), billing and invoicing automation (used by 32 percent) and client information management (used by 28 percent).[195] These are the foundational transactional functions where small businesses process the highest volumes of repetitive, rule-governed work, and where the data is cleanest, the processes are most defined and the outcomes are most directly measurable.

This matters enormously for where agents land first. An AI agent deployed into a workflow that is already partially digitized inherits clean-enough data to operate on, a defined process to execute against and a baseline performance metric against which its contribution can be measured. An agent deployed into a workflow that is still primarily manual inherits the opposite: inconsistent data, undefined process and no baseline. The operational functions where SMBs have already concentrated their digital tool usage are precisely the functions where agent deployment is most likely to succeed on a short timeline; and the functions where the case for deployment is easiest to make because the manual labor being replaced is already visible, quantifiable and universally recognized as a friction point.

THE ASPIRATION-REALITY GAP

The behavioral data matters particularly because it sits in tension with what SMBs say they want AI to do when asked directly. Salesforce's global survey of SMB leaders identifies the prominent aspired AI use cases as marketing campaign optimization, content generation and customer experience personalization.[196] These are genuinely valuable applications. They are also applications that require integrated customer data, clean CRM records and workflow architecture that most small businesses have not yet built. The gap between the marketing optimization aspiration and the payment collection reality is a gap in data readiness and operational maturity; the same gap that Chapter 8 documented as the single most common proximate cause of AI deployments that underdeliver.

The operations-first adoption pattern is therefore not a limitation to be overcome on the path to more sophisticated AI use. It is the correct sequence. The business that successfully deploys an agent to handle invoice follow-up has, in doing so, proved that its billing data is clean enough to act on autonomously, its client communication workflow is defined enough to execute reliably and its team can manage a human-agent handoff effectively. Each of those proofs is a prerequisite for the more complex, customer-facing AI deployments that come later. The path from operational agent to strategic agent runs directly through the back-office stack, and the providers who understand this sequence are the ones with successful deployments rather than stalling out.

195 vcita. "Survey Results: What Do SMBs Want in 2025?" inTandem by vcita, Mar. 2025.

196 Salesforce. "Small and Medium Business Trends Report, 6th Edition." Salesforce, 2025.

THE DIGITAL LABOR STACK AND THE SMB FIT

Recent research on the agentic labor economy documents the organizational dimension of the operations-first pattern at the leadership level. 67 percent of CEOs believe that implementing AI agents is critical for their organization to remain competitive, and 73 percent agree that agentic labor will transform their company structure.[197] But the analysis also identifies the functions where agentic AI augmentation is expected to have the greatest near-term impact (sales operations, marketing operations and IT operations) all of which, in the SMB context, map directly to the back-office and operational workflows where digital tool adoption has already concentrated.[198]

The operational functions that dominate SMB digital tool usage (payment collection, invoicing, client management, scheduling) are not just where small businesses have their cleanest data. They are also the functions that consume the most time relative to the strategic value they generate. A small business owner who spends 90 minutes each morning processing invoices, following up on late payments, updating client records and reconciling accounts is spending 90 minutes on work that is entirely rule-governed and repeatable, and ideal for autonomous execution. That is the work that digital labor is structurally best suited to absorb; and it is the work that, when absorbed, frees the human capacity that the aspired customer experience and marketing AI use cases require.

THE EMBEDDED DELIVERY MODEL

The vcita/inTandem data also illuminates the delivery model question that determines whether digital labor adoption happens or gets deferred indefinitely. In that survey, ease of use ranked as the primary factor in SMB technology purchasing decisions; consistently outranking price, feature depth and vendor reputation.[199] Among SMBs that reported hesitancy about AI automation, the dominant concern was not cost or security but unfamiliarity: they wanted to experience AI before permitting it to take autonomous action on their behalf.[200]

This behavioral reality has a direct structural implication. Agentic labor in the SMB market will not primarily arrive through standalone AI agent platforms that require dedicated configuration, integration and onboarding. It will arrive through the business management tools that small businesses already use, embedded in the billing systems, CRM platforms, payment processors and scheduling tools where daily work already happens. An AI capability that surfaces inside a tool a business owner uses every morning requires no new login, no new interface, no new workflow habit. It reduces the activation energy of adoption to near zero. A standalone agent platform that requires the business owner to learn a new system, migrate data and redesign their workflow before seeing any value faces exactly the activation barriers that explain why 92 percent of SMBs have adopted digital tools but only 10 percent have fully integrated them.

For technology partners, this has a clear strategic implication: the highest-conversion agentic labor conversation is not "here is a new AI platform I would like you to adopt." It is "here are the AI agent capabilities already embedded in the tools you are paying for today, and here is how we activate them." The partners who have mapped their clients' existing software stacks against the AI capabilities available within those stacks, and who show up with an activation roadmap rather than a new vendor proposal, are the partners who turn the operations-first adoption pattern into recurring revenue.

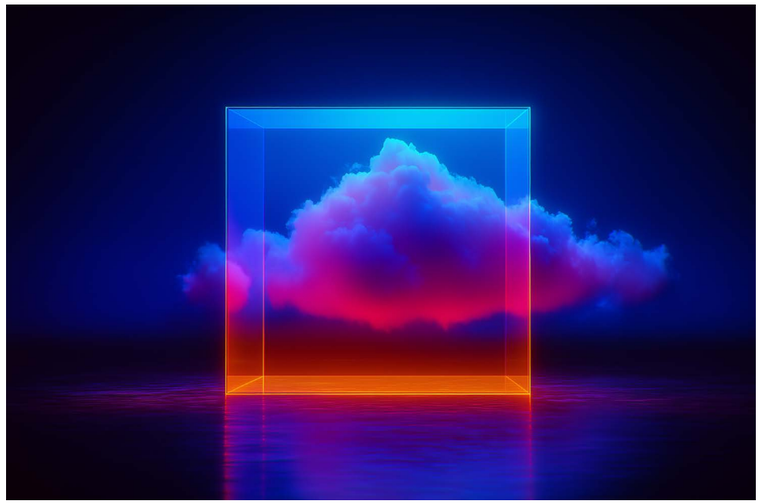

THE BEACHHEAD AND THE EXPANSION PATH

The operations-first pattern is not the destination. It is the beachhead. A small business that has successfully deployed an agent to handle payment follow-up has not completed a digital labor strategy, but the first proof; the demonstration that autonomous execution is reliable and valuable in a bounded context; that makes the next deployment easier to justify and faster to implement.

> The operations-first pattern is not the destination.

IDC's 2026 SMB Digital Landscape research identifies the direction of that expansion. As SMBs progress from basic AI tool adoption to strategic deployment, the use cases that generate the highest perceived value shift from operational automation toward customer engagement, sales acceleration and competitive differentiation.[201] Marketing and customer experience AI, the use cases SMBs say they want but are not yet ready for, are not the starting point. They are the destination that the operational beachhead makes reachable.

The managed intelligence provider who sequences this journey deliberately; starting with the operational functions where data is cleanest and success is most predictable, building the data infrastructure and governance frameworks required for more complex deployments along the way, and expanding the agent footprint as each proof compounds into the next; are building the institutional knowledge of each client's operational architecture that makes every subsequent deployment faster, more accurate and more defensible. The switching cost that compounds with each new phase of the journey is the natural consequence of deep, sequenced expertise that no competitor can replicate from a standing start.

The endpoint of that expansion path is worth naming explicitly, because it reframes the managed intelligence opportunity from a sequence of deployments into a workforce architecture question. Over time, the businesses that have completed the journey from operational beachhead to strategic deployment are likely to arrive at a recognizable organizational model: every employee paired with an always-on agent that functions as the front-end operating layer for their work. That agent will not simply answer questions on demand. It will monitor open commitments, surface next actions, coordinate downstream systems, request approvals and increasingly prompt the human rather than wait to be prompted, managing a broader network of task-specific agents behind the scenes on the employee's behalf. The visible relationship may appear to be one employee and one agent. Beneath the surface, that personal agent may be orchestrating dozens of specialized agentic tasks across billing, scheduling, compliance, client communication and internal knowledge management simultaneously. The SMB that reaches this architecture has assembled a digital workforce and the technology partner who sequenced that journey, from the first operational agent to the persistent employee-facing layer, is the partner whose institutional knowledge of that business is by then inseparable from the way the business runs.

197 IDC. "Voice of the CEO on Digital Labor." IDC InfoBrief, sponsored by Salesforce, Oct. 2025.

198 IDC. "Voice of the CEO on Digital Labor." IDC InfoBrief, sponsored by Salesforce, Oct. 2025.

199 vcita. "Survey Results: What Do SMBs Want in 2025?" inTandem by vcita, Mar. 2025.

200 vcita. "Survey Results: What Do SMBs Want in 2025?" inTandem by vcita, Mar. 2025.

201 Evans, Katie, et al. "IDC FutureScape: Worldwide Small and Medium-Sized Business 2025 Predictions." IDC, 2024.

THE KNOWLEDGE POLLUTION PROBLEM

AI abundance will not arrive only as more productivity. It will arrive as more volume: more knowledge, more artifacts, more communication, more workflows and more machine-generated demands for human attention.

The first wave of AI made this visible without solving it. Consider what a knowledge worker does when they interact with an AI tool today: they see a meeting transcript, decide whether it matters, copy it into the AI interface, explain the context, craft the prompt, evaluate the output and decide what to do next. Every step after the first is work the human created for themselves by using the tool. The employee became the go-between, the connective tissue between AI capability and business outcome. That is a transitional burden and one that runs in both directions simultaneously.

Internally, AI-generated knowledge compounds: more artifacts, more summaries, more drafts, more workflow outputs, more machine-generated requests for human review and approval. Externally, every vendor, partner and competitor is now deploying outbound agents optimized to capture human attention at scale; filling inboxes, calendars, CRMs, support queues and Slack channels with AI-generated communication that competes for the same finite resource. One organization's AI growth engine is another employee's AI-generated noise.

The knowledge worker is caught in an arms race for their own attention; one they cannot win by working harder, moving faster or adding more tools to the stack. The only answer is a new layer of architecture.

THE ALWAYS-ON AGENT: THE FILTER AND THE FIREWALL

Every employee, from day one, is paired with an always-on agent scoped to their identity and their role. It knows their calendar, their context, their history. Context windows already span millions of tokens and soon, they will hold the entire operational memory of a business. That agent knows the meeting from six months ago, the decision that shaped last quarter, the client relationship that goes back three years.

The first phase of AI was humans prompting machines. The next phase is machines prompting humans. The personal persistent agent is both the filter for everything coming in and the firewall against everything that should not reach the human at all. In fact, sixty-seven percent of executives already agree that AI agents will drastically transform existing roles, not by eliminating them, but by restructuring them around the always-on agent as the primary operating interface.[202]

The moment this stopped being theoretical has a precise date. In November 2025, Austrian developer Peter Steinberger published an open-source autonomous agent framework; initially named Clawdbot, later renamed OpenClaw; designed to run persistently on a user's own devices, integrating across messaging platforms, monitoring open loops and executing tasks autonomously without session-based prompting.[203]

OpenClaw surpassed 250,000 GitHub stars in approximately 60 days, a milestone that took React more than a decade to reach, making it one of the fastest-growing open-source projects in history.[204] Before OpenClaw, AI waited to be asked. OpenClaw changed the nature of the relationship: for the first time, AI was persistent, watching, listening and acting without being asked. IBM Research observed that OpenClaw demonstrated the real-world utility of always-on agents is not limited to large enterprises; a finding reflected in the rapid growth of the ClawHub skill marketplace, which crossed 3,000 community-built integrations within weeks of launch, spanning marketing automation, CRM connectivity, and productivity workflows.[205]

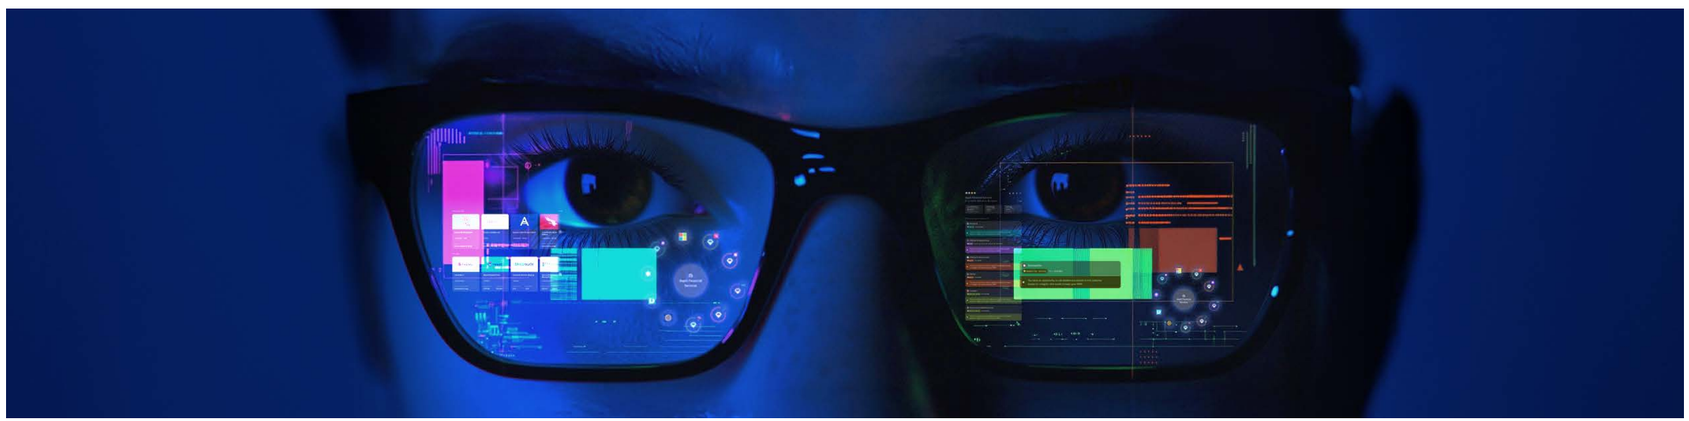

THE TASK-BASED AGENT: THE EXECUTION LAYER

The second type of agent operates at the business function level. Every core workflow of the business gets its own purpose-built, governed, accountable agent. As companies transform legacy processes and human-oriented workflows with AI, these task-based agents become the new execution layer of the business; handling recurring work across sales, finance, marketing, support, operations, HR, legal, procurement and analytics.[206] Over 70 percent of AI rollout initiatives now focus on action-based agents rather than conversational assistants, and among organizations that have adopted them, nearly two-thirds report increased productivity, more than half report cost savings and 55 percent report faster decision-making.[207] These are the agents that land first for precisely the reasons established earlier in this chapter: they operate where digital infrastructure already exists, data is cleanest and the outcome baseline against which their performance is judged is already visible. The task-based agent is the beachhead unit of the digital workforce.

HOW THE TWO TYPES CONNECT

The always-on personal agent becomes the employee's representative inside the entire agentic infrastructure. It delegates to task agents. Coordinates across domains. Monitors outcomes. Escalates only what requires human intervention.

The traditional organizational hierarchy ran from executive to manager to individual contributor; layer upon layer of coordination, with humans creating work for other humans. The intelligence-native business works differently. Humans orchestrate. Personal agents direct. Task-based agents execute. Every employee becomes a manager, not of people, but of a portfolio of agents working on their behalf. The two types of agents are complementary layers of the same architecture. The always-on personal agent becomes the employee's interface to the full stack of task-based agents running beneath it; delegating to them, monitoring their outputs, escalating only what requires human intervention and coordinating across domains on the employee's behalf. Organizations with a named agent owner, someone accountable for overseeing that coordination, have a 2.7 times higher production conversion rate than those without one,[208] a finding that confirms what the architecture implies: the value of the two-agent model is realized by governing the relationship between them. Over time, what appears from the outside to be one employee working alongside one agent may, beneath the surface, represent a single personal agent orchestrating dozens of simultaneous task-level executions across every function the employee touches.

That is the formation of the agentic workforce: not humans replaced by AI, and not humans manually prompting AI, but humans augmented by always-on agents that orchestrate networks of specialized agentic labor.[209] The visible relationship may appear to be one employee working alongside one agent. Beneath the surface, that personal agent may be coordinating dozens of simultaneous task-level executions across every function the employee touches. The SMB that reaches this architecture has assembled a digital workforce, and every agent across both types must be governed, monitored and accountable. That governance requirement is the foundation everything else is built on, and it is the subject of Chapter 16.

202 PwC. "AI Agent Survey: From Adoption to Transformation." PwC, May 2025.

203 NVIDIA Blog. "What OpenClaw Agents Mean for Every Organization." NVIDIA, Apr. 2026.

204 NVIDIA Blog. "What OpenClaw Agents Mean for Every Organization." NVIDIA, Apr. 2026.

205 Wikipedia. "OpenClaw." Wikipedia, last modified Apr. 2026.

206 PwC. "AI Agent Survey: From Adoption to Transformation." PwC, May 2025.

207 PwC. "AI Agent Survey: From Adoption to Transformation." PwC, May 2025.

208 Digital Applied. "AI Agent Adoption 2026: 120+ Enterprise Data Points." Digital Applied, Apr. 2026.

209 PwC. "AI Agent Survey: From Adoption to Transformation." PwC, May 2025.

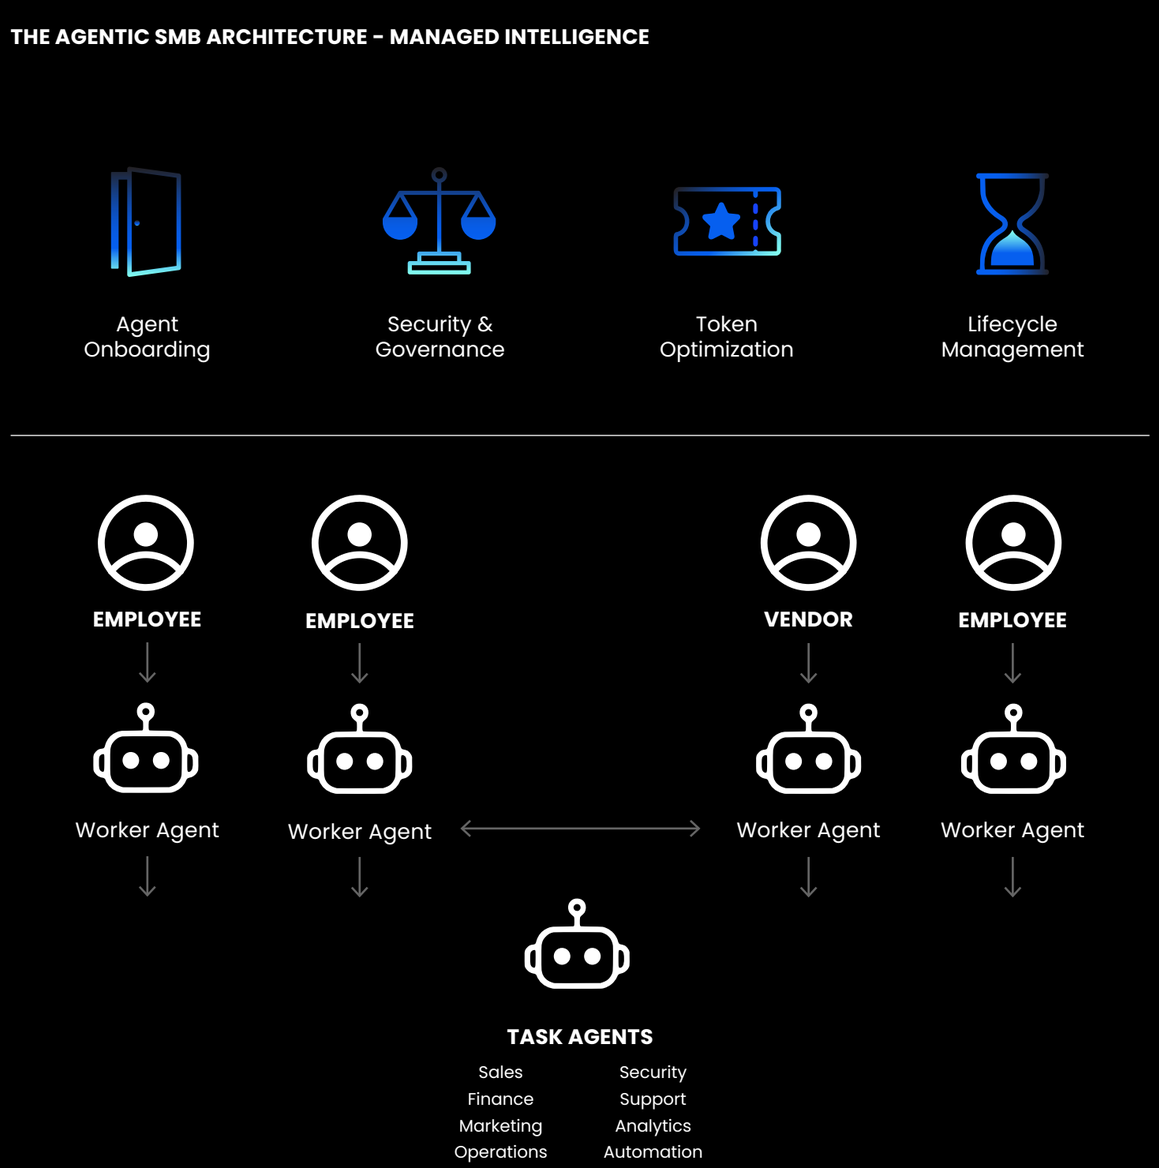

HUMAN CAPITAL + TOKEN CAPITAL

- Agentic workflows turn intelligence into a **managed resource** - much like headcount
- **Proactive always-on** worker agents will be assigned to every employee
- Vendor's will also **forward-deploy** persistent worker agents
- **Task agents** will emerge as orchestrated fleets of short-lived capabilities, **managed by both human and agent workers**

Agentic workflows turn intelligence into a managed resource - much like headcount.

Chapter 15 examines the vertical dimension of this same argument: the finding that the sequenced deployment strategy that works across the SMB economy works fastest and stickiest when it is organized around the specific workflows, compliance requirements and business rhythms of a single industry rather than applied generically across all of them.

Chapter 15.

Vertical Playbooks Over Horizontal Platforms

There is a pattern in every technology adoption cycle that the managed services channel has lived through; a period of horizontal generalism followed by a period of vertical specialization. The first phase of the internet economy produced general-purpose web developers. The second produced specialists in e-commerce, healthcare portals, financial services applications and others. The first phase of cloud migration produced general-purpose managed service providers. The second produced specialists in healthcare cloud compliance, financial services security architecture and manufacturing operations monitoring. The AI adoption cycle is following the same pattern, and the channel is currently at the inflection point where the generalist phase is ending and the specialist phase is beginning.

The evidence for this inflection is both structural and empirical. It is structural in that the problem of AI adoption in the SMB economy is not a problem that a generic solution can solve. The workflows that need to be redesigned are industry specific. The compliance requirements that govern what agents can and cannot do autonomously are industry specific. The data formats, terminology, role structures and operational rhythms that AI systems must understand to operate effectively are industry specific. A horizontal AI platform that works adequately across all industries works excellently in none of them; and in a market where 70 percent of SMBs will demand documented use cases and measurable outcomes before committing AI investment, "adequately across all industries" is not a proposition that converts.[210]

It is empirical in that the vertical SaaS market (the closest available proxy for the vertical AI services market the channel is now entering) is growing at two to three times the rate of horizontal platforms, driven by precisely the dynamics that make vertical specialization structurally advantageous in an AI-agent world.[211]

While 79 percent of organizations are experimenting with generative AI, fewer than 10 percent have successfully scaled AI agents in production.

WHY HORIZONTAL STALLS

The Omdia research on MSP channel dynamics documented in Chapter 10 captures the production failure of horizontal AI deployment: 61 percent of channel partners currently struggle to move AI projects beyond proof-of-concept with their existing clients.[212]

The failure mode that Omdia describes is one of context. The partner lacks the workflow redesign expertise, the domain-specific data architecture knowledge and the industry-specific governance frameworks required to move a demonstrated capability from a controlled demonstration environment into the messy, regulation-constrained, legacy-data reality of an actual business operating in an actual industry.

McKinsey's April 2026 analysis of agentic AI deployment confirms the scale of this stall at the broader market level: while 79 percent of organizations are experimenting with generative AI, fewer than 10 percent have successfully scaled AI agents in production.[213] The gap between 79 percent experimenting and 10 percent scaled is pilot purgatory; the condition in which AI investment generates demonstrations without generating returns, and in which the partner who cannot guide the client from proof to production becomes a liability rather than an asset.

The root cause of pilot purgatory is almost always the same. A horizontal AI implementation begins with a compelling demonstration of what the technology can do in a generalized context. It then encounters the specific: the legal firm's intake workflow includes exception handling for jurisdictional conflicts that the general agent configuration does not anticipate. The medical billing follow-up agent runs into prior authorization requirements that vary by payer in ways the horizontal platform has no pre-built logic to accommodate. The construction compliance tracking deployment collides with permit databases that are formatted differently across every municipality in the client's operating geography. Each of these is a solvable problem, but solving it requires domain expertise that a horizontal provider cannot credibly claim and cannot efficiently build from scratch for every client in every industry they serve.

210 Evans, Katie, et al. "IDC FutureScape: Worldwide Small and Medium-Sized Business 2025 Predictions." IDC, 2024.

211 Lemkin, Jason M. "Why SaaS Companies That Sell Outside of Tech Are on Fire." SaaStr, 1 Jul. 2025. See also: Modall. "SaaS Trends 2025–2026: 25 Definitive Trends Shaping the Industry." Modall, Dec. 2025.

212 Omdia. "MSP Trends and Predictions 2025 — Executive Summary." Omdia (Canalys), Jan. 2025.

213 Singla, Alex, et al. "The State of AI in 2025: Agents, Innovation, and Transformation." McKinsey & Company, Nov. 2025. Additionally: McKinsey & Company. "McKinsey and Wonderful Team Up to Deliver Enterprise AI Transformation from Strategy to Scale."

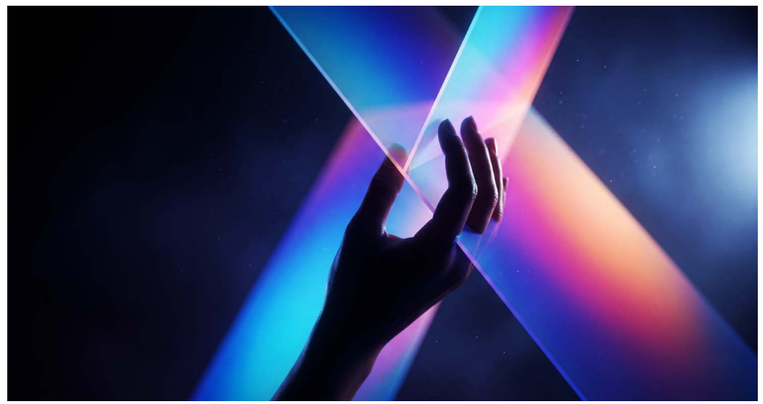

WHY VERTICAL WINS
The vertical SaaS growth premium is the market's running measurement of what domain knowledge is worth; the documented price difference between software built for an industry and software adapted to survive in one. SaaS platforms are growing at two to three times the rate of horizontal counterparts, with net revenue retention rates 20 to 30 percent higher and churn rates up to 50 percent lower; structural performance advantages that flow directly from the depth of workflow integration that vertical specialization enables.[214]

For the MIP, the implication is direct. The partner who has built deep expertise in healthcare, legal services, construction or home services does not start each new AI engagement from zero. They start with a workflow map, a compliance framework, a data architecture template and a set of outcome metrics that they have refined across multiple deployments in the same industry. The second deployment is faster than the first. The third is faster than the second. Each engagement compounds the domain knowledge that makes the next one more defensible.

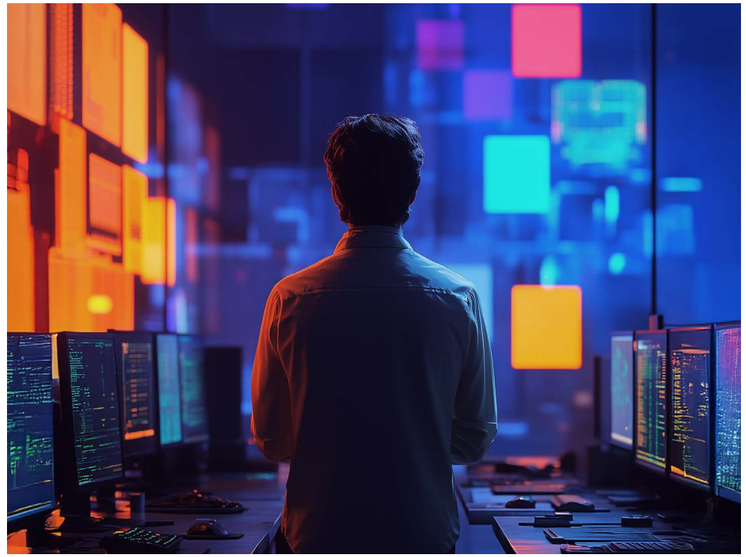

SaaS platforms are growing at two to three times the rate of horizontal counterparts.

The mechanism is the same one that will drive vertical AI services to outperform horizontal AI services in the SMB market: a solution built for a specific industry does not require the client to explain their context. The dental practice management platform already knows what a treatment plan approval workflow looks like. The construction management platform already knows what a lien waiver tracking process requires. The professional services billing platform already knows what timekeeper rate management involves. When AI agents are deployed within these contexts, they inherit domain knowledge that horizontal platforms must build from scratch; and the difference in deployment speed, error rate and client confidence is measurable.

THE PLAYBOOK ARCHITECTURE
The practical expression of vertical specialization is a playbook: a structured engagement model that packages the domain expertise, deployment sequence, governance architecture and outcome metrics for a specific high-value workflow in a specific industry into a repeatable, fixed-scope offering that can be delivered consistently across a portfolio of clients in that vertical.

The key insight is that the playbook solves the proof-of-concept problem by clearly defining the path from demonstration to production. When a managed intelligence provider deploys a legal intake triage agent, they are not demonstrating that AI can handle intake triage. They are delivering a pre-configured, pre-governed, pre-measured system that handles intake triage for a law firm; with the specific exception categories, escalation triggers, conflict check integrations and billing code classifications that law firm intake requires already built in. The client does not need to define the scope, the guardrails or the success metrics. The playbook has already defined them and expressed as a documented deliverable with a 90-day implementation timeline and measurable outcome commitments.

IDC's research on SMB technology buying behavior documents why this matters so acutely in the current market. The 70 percent of SMBs who will demand clear, documented use cases before committing new AI investment are asking for a business case, proof that a specific capability, delivered in a defined scope, has produced measurable results for a business like theirs.[215] The vertical playbook is the business case made operational. It is a deployment framework that doubles as a sales asset, because the documented outcomes from prior deployments in the same industry are the most credible possible evidence that the next deployment will succeed.

WHERE PLAYBOOKS ARE BUILT: THREE INDUSTRY ARCHETYPES
The industries where vertical AI playbooks generate the fastest return on investment share three characteristics: high workflow volume, strong regulatory pressure and established digital tool usage in the operational functions where agents land first. Three archetypes illustrate the pattern.

Professional and legal services present perhaps the clearest foothold. Intake triage: the classification, routing and initial processing of new client inquiries is a high-volume, rule-governed workflow that consumes significant attorney and paralegal time without requiring the judgment and relationship skills that define professional value. A legal intake triage agent that classifies inquiry type, checks for conflicts, assigns to practice area, generates an initial acknowledgment and populates the CRM record is executing work that is entirely repeatable and appropriate for autonomous handling. The compliance dimension (bar association communication requirements, conflict disclosure obligations, privilege preservation) is industry-specific but consistent enough across practices to be encoded in a playbook rather than rebuilt from scratch for each client.

214 Lemkin, Jason M. "Why SaaS Companies That Sell Outside of Tech Are on Fire." SaaStr, 1 Jul. 2025. See also: Modall. "SaaS Trends 2025–2026: 25 Definitive Trends Shaping the Industry." Modall, Dec. 2025.

215 Evans, Katie, et al. "IDC FutureScape: Worldwide Small and Medium-Sized Business 2025 Predictions." IDC, 2024.

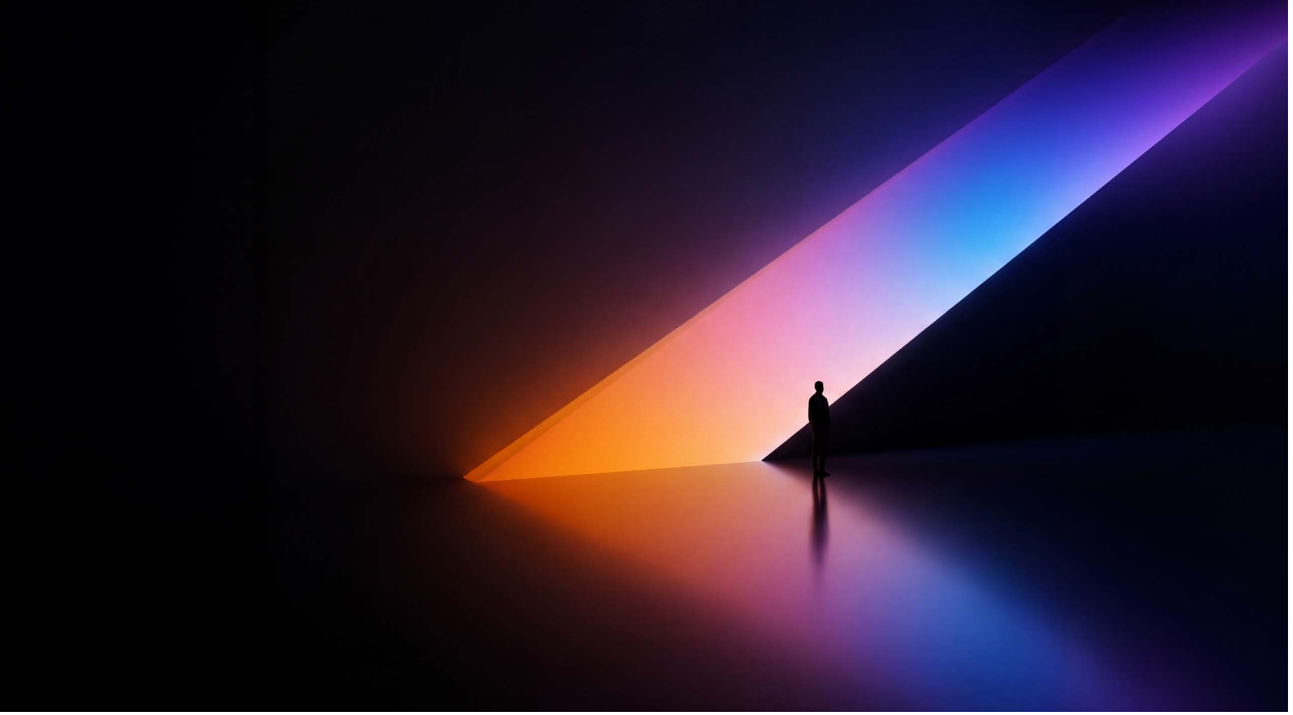

A managed intelligence provider who deploys their first legal intake triage playbook learns something that their second deployment refines.

Healthcare and wellness practices present a parallel opportunity in billing follow-up and prior authorization management. Medical billing follow-up is among the highest-volume, lowest-judgment workflows in any healthcare practice, and it is also among the most consequential for cash flow. An agent that monitors outstanding claims, identifies claims approaching denial deadlines, generates follow-up communications to payers and escalates complex cases to billing staff is handling work that practices currently manage manually with significant administrative overhead. The payer-specific knowledge required to make this agent effective (prior authorization requirements, appeal windows, documentation standards) is industry-specific domain expertise that a vertical playbook encodes and that a horizontal platform cannot pre-populate.

Construction and trades present a governance-heavy entry point in compliance tracking and permit management. Construction projects generate compliance obligations across multiple jurisdictions, multiple trade categories and multiple regulatory frameworks simultaneously. An agent that tracks permit status, generates inspection request communications, monitors lien waiver deadlines and flags compliance gaps before they become project delays is executing work that project managers currently track in spreadsheets and calendars, with predictable gaps, delays and cost consequences. The jurisdiction-specific knowledge that makes this agent effective is the domain expertise that a vertical construction practice has already accumulated through years of project management in specific geographies.

THE COMPETITIVE MOAT THAT PAYS OFF

The vertical playbook strategy creates a competitive position that horizontal generalism cannot replicate from a standing start, and the mechanism of that advantage is time-dependent in a way that creates urgency for providers who have not yet chosen their verticals.

A managed intelligence provider who deploys their first legal intake triage playbook learns something that their second deployment refines. The second deployment teaches them something the third deployment benefits from. Each deployment adds to the institutional knowledge of how law firms operate, what the common exception cases are, where the compliance boundaries sit and what outcome metrics clients care about enough to renew a managed service contract around. After ten deployments in legal services, the provider has a body of operational knowledge; encoded in their playbook, configuration templates, governance frameworks and outcome benchmarks; that a competitor entering the vertical from a horizontal generalist position would require years of deployments to replicate.

The vertical SaaS analogy is instructive here. Procore did not become the dominant construction management platform by having better technology than generic project management tools. It became dominant by knowing more about how construction projects work than any horizontal platform could learn from serving thousands of industries simultaneously. Its data moat, its workflow depth and its compliance pre-configuration represent years of domain accumulation that no competitor can shortcut. The MIP who builds the equivalent depth in two or three verticals is building the same kind of structural advantage; one that compounds with each deployment and becomes progressively more difficult to displace as clients' operations become more deeply integrated with a partner who understands their industry from the inside.

Chapter 16 examines the final and most overlooked revenue opportunity in the managed intelligence portfolio: the AI governance audit, the discovery and remediation of the shadow AI that most SMB clients are already running without their technology partner's knowledge and why the audit conversation is the highest-margin, highest-retention entry point available to providers who know how to lead it.

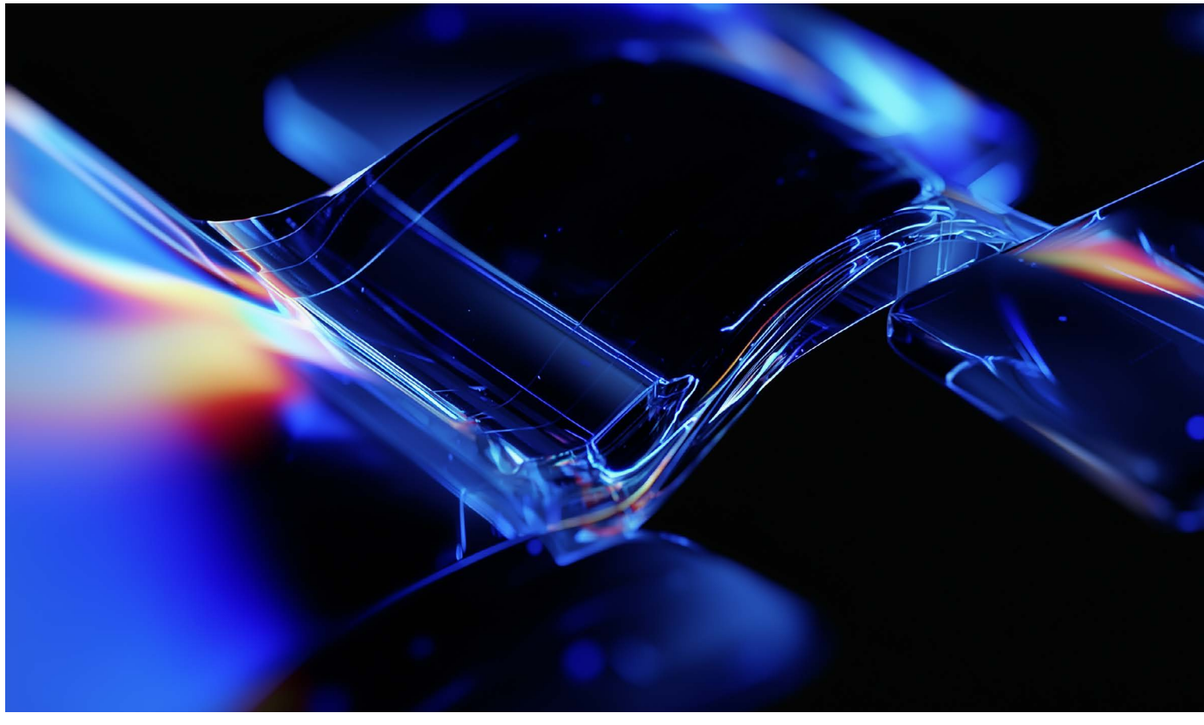

Chapter 16.

The AI Governance Audit: The Most Overlooked Revenue Line

The managed intelligence conversation almost always begins with deployment.

- Which workflows should be automated first?
- Which agents should be configured and in what sequence?
- What the data architecture needs to look like before the first autonomous action is taken?

These are the right questions, but they are the right questions for a client who is starting from zero, the AI-native. Most SMB clients are not starting from zero. They are starting from somewhere in the middle, with AI tools already running in their environment that their technology partner does not know about, has not evaluated and has no visibility into. The governance audit is the conversation that begins there, and it is, in both economic and strategic terms, the most underutilized entry point in the managed intelligence portfolio.

The reason it is underutilized is because the dominant framing of the market opportunity positions the technology partner as the gateway through which AI arrives in the small business economy. That framing is wrong. The U.S. Chamber of Commerce's 2025 technology survey found that 58 percent of SMBs have already adopted AI in some form, up from 40 percent just one year earlier.[216] The gateway has already opened. What has not followed through on is governance, policy or oversight.

ConnectWise's cybersecurity research finds that 49 percent of SMBs do not have AI-specific security policies in place.[217] The providers who understand this are the same ones addressing a governance emergency that is already underway in their existing client base.

AI tools are freely available and arrive in the workflow before any governance framework has been built to receive them.

THE SHADOW AI POPULATION IS LARGER THAN ANYONE IS MEASURING

The WalkMe 2025 AI in the Workplace Survey found that 78 percent report using unapproved AI tools at work, with only 7.5 percent receiving anything resembling extensive AI training from their employers.[218] Gartner's survey of cybersecurity leaders confirms the organizational dimension: 69 percent of organizations have evidence or strong suspicion that employees are using prohibited generative AI tools.[219] Microsoft's research finds that 71 percent of UK workers admit to using unapproved AI tools at work, with 22 percent applying them to critical financial tasks.[220]

The broader pattern of governance lagging capability is documented independently by the Stanford AI Index, which tracks reported AI incidents across industries. Documented incidents reached 362 in 2025, up from 233 the year before; a rate of escalation that reflects ungoverned AI deployment, where the gap between what systems are capable of doing and what policies have been established to contain them continues to widen.[221]

These cases are examples of the behavioral norm in a market where AI tools are freely available and arrive in the workflow before any governance framework has been built to receive them. The employee who discovers that a consumer AI tool summarizes client emails faster than their approved workflow is not circumventing policy out of malice. They are solving a productivity problem with the most accessible available tool, and in doing so they are introducing company data into AI systems that the business has not evaluated, approved or configured to meet its compliance and security obligations.

Harmonic Security's analysis of 22 million enterprise prompts from 2025 documents what that data looks like: code, legal documents and financial data comprise 74.5 percent of what employees expose through unsanctioned AI tools, with legal documents (merger and acquisition materials, settlement content, litigation strategy) representing the single largest category at 35 percent.[222]
The shadow AI problem is not employees using AI for low-stakes drafting tasks. It is employees routing the most sensitive information in the business through AI systems their organization has never reviewed.

THE FINANCIAL CONSEQUENCES ARE NO LONGER THEORETICAL

IBM's 2025 Cost of Data Breach Report establishes the cost dimension that should anchor every governance audit conversation. Shadow AI incidents now carry a cost premium over standard breaches: $4.63 million per incident compared to $3.96 million for standard breaches, a differential of more than $650,000 per event.[223] The same research finds that 97 percent of AI-related breaches lacked proper access controls; meaning the financial exposure is not distributed randomly across the population of AI-deploying organizations. It is concentrated in the organizations that deployed AI without the access governance architecture that makes autonomous action safe.

For an SMB whose annual revenue may fall between $1 million and $10 million, a $4.63 million breach is not a manageable cost event. The governance audit conversation then becomes a survival argument, and it is one the provider can make with primary data, not hypothetical risk framing.

216 U.S. Chamber of Commerce. "Empowering Small Business: The Impact of Technology on U.S. Small Business." U.S. Chamber of Commerce, 13 Aug. 2025.

217 ConnectWise. "SMB Cybersecurity Statistics and Trends." ConnectWise, 2025.

218 WalkMe. "New WalkMe Survey Shows Shadow AI Is Rampant; Training Gaps Undermine AI ROI." SAP/WalkMe, Aug. 2025.

219 Gartner. "Gartner Predicts 40% of AI Data Breaches Will Arise from Cross-Border GenAI Misuse by 2027." Gartner, 17 Feb. 2025.

220 Microsoft. "Rise in 'Shadow AI' Tools Raising Security Concerns for UK." Microsoft UK Stories, 13 Oct. 2025.

221 Sajadieh, Sha, et al. "The AI Index 2026 Annual Report." AI Index Steering Committee, Institute for Human-Centered AI, Stanford University, Apr. 2026.

222 Harmonic Security. "What 22 Million Enterprise AI Prompts Reveal About Shadow AI in 2025." Harmonic Security, Jan. 2026.

223 IBM Security. "Cost of a Data Breach Report 2025." IBM, 2025.

THE AGENT ACCELERATION MAKES THIS URGENT NOW
The governance problem that currently exists at the level of ungoverned SaaS-based AI tools is about to become structurally more complex. Gartner's August 2025 research shows the pace: fewer than 5 percent of enterprise applications are currently integrated with task-specific AI agents, a figure Gartner predicts will reach 40 percent by the end of 2026.[224] The IDC and Microsoft research on the agent economy projects 1.3 billion agents in deployment by 2028.[225]

The governance gap that exists today around passive AI tools, tools that generate outputs employees act on, will be qualitatively different when it exists around agentic AI systems that take actions autonomously. An ungoverned AI tool that drafts a client email creates a review opportunity before the email is sent. An ungoverned AI agent with send permissions creates no such opportunity. The audit that discovers and governs passive shadow AI today is the same audit that establishes the governance architecture for agentic deployment tomorrow, and the technology provider who has completed that audit owns the only credible position from which to guide the client's agentic transition safely.

Deloitte's State of AI 2026 frames the organizational reality that makes this window both urgent and finite: close to three-quarters of organizations plan to deploy agentic AI within two years, but only 21 percent currently report a mature model for governing autonomous agents.[226] The 79 percent without mature governance are not planning to remain ungoverned. They are planning to deploy agents into governance frameworks that do not yet exist, on timelines that are already running. The technology provider who arrives before those deployments begin, with an audit that maps the current AI environment and a governance framework that makes agentic deployment safe, is not selling a service. They are positioning themselves as the prerequisite for everything that follows.

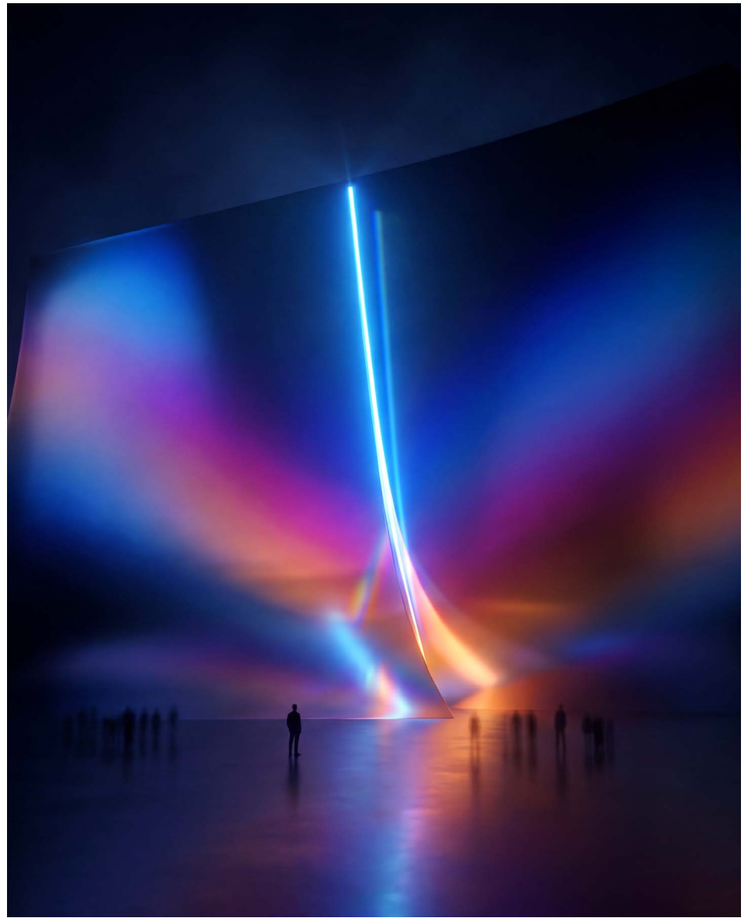

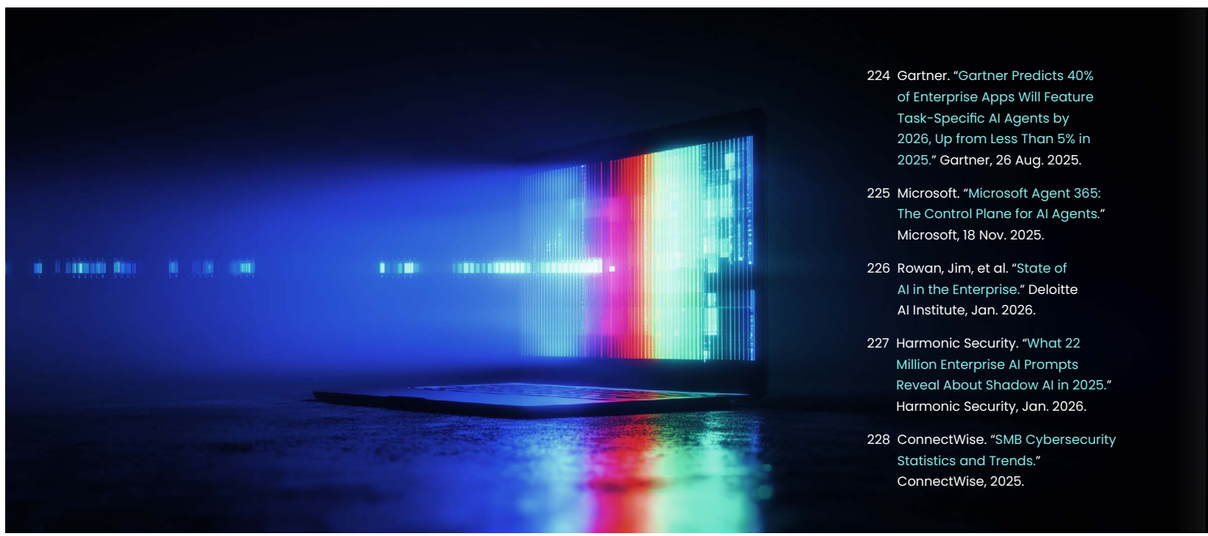

224 Gartner. "Gartner Predicts 40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026, Up from Less Than 5% in 2025." Gartner, 26 Aug. 2025.

225 Microsoft. "Microsoft Agent 365: The Control Plane for AI Agents." Microsoft, 18 Nov. 2025.

226 Rowan, Jim, et al. "State of AI in the Enterprise." Deloitte AI Institute, Jan. 2026.

227 Harmonic Security. "What 22 Million Enterprise AI Prompts Reveal About Shadow AI in 2025." Harmonic Security, Jan. 2026.

228 ConnectWise. "SMB Cybersecurity Statistics and Trends." ConnectWise, 2025.

The AI Governance Audit Framework

1. Inventory
2. Risk Assessment
3. Remediation
4. Policy

The governance gap that exists today around passive AI tools, tools that generate outputs employees act on, will be qualitatively different when it exists around agentic AI systems that take actions autonomously.

WHAT THE AUDIT DISCOVERS
The AI governance audit is a structured discovery process with four components that map directly to billable remediation work.

The first component is inventory. This is a systematic identification of every AI tool currently in use across the client's environment, whether procured through official channels, adopted by individual employees using personal accounts or embedded within existing SaaS platforms the client already pays for. Most SMBs have no comprehensive inventory. The audit creates one, and in doing so makes visible a population of tools, data flows and autonomous actions that the business has been operating without knowing.

The second component is risk assessment. This is an evaluation of each identified tool against the client's data classification requirements, compliance obligations and security architecture. Harmonic Security's data makes the stakes concrete: the average organization experiences 223 incidents per month of users sending sensitive data to AI applications, double the rate from the prior year.[227] The risk assessment translates that aggregate exposure into client-specific findings; which tools are handling regulated data, which are processing information that triggers compliance obligations, which are operating with access credentials that create lateral movement risk in the event of compromise.

The third component is remediation. This is where the audit converts from a discovery engagement into a managed service. Ungoverned tools are either retired, reconfigured with appropriate data handling controls or integrated into the client's managed AI stack with monitoring, access governance and audit trail architecture. For the technology provider, this is the highest-margin work in the digital labor portfolio because it requires the deepest institutional knowledge of the client's environment, creates the strongest switching costs of any engagement type and generates the ongoing monitoring relationship that is the structural foundation of managed intelligence.

The fourth component is policy. This is the development of an AI acceptable use framework that defines which tools are approved, how data may be handled within them, what actions agents may take autonomously and under what conditions and what the escalation and incident response process looks like when an AI system behaves outside its defined parameters. ConnectWise's research finding that 49 percent of SMBs lack AI-specific security policies means that for roughly half of any managed services portfolio, the policy component of the audit is net-new work with no prior framework to build on.[228] That signifies legitimate scope expansion.

The agentic AI monitoring and observability tools market stands at $0.55 billion in 2025 and is forecast to reach $2.05 billion by 2030.

Underlying all four audit components, and the ongoing monitoring relationship that follows, is a capability that the agentic enterprise cannot govern without: observability. AI observability is the practice of collecting data from each AI action to enable system transparency and understandability, so organizations can see not just what is happening, but why.[229] In a traditional software environment, testing can be comprehensive because inputs are finite and predictable. In an agentic environment, that certainty is not available: agents make autonomous decisions, interact with external tools, process unstructured data and generate outputs that vary even with identical inputs. This non-deterministic nature introduces unique challenges including hallucinations, performance drift, unexpected behaviors and compliance violations that traditional monitoring approaches cannot adequately address.[230] The observability layer: logs, traces, model outputs, data flows and decision paths; is what converts an agentic deployment from a governance liability into a governed asset.

The agentic AI monitoring and observability tools market stands at $0.55 billion in 2025 and is forecast to reach $2.05 billion by 2030, with SMBs registering the fastest growth rate of any segment at 32.4% CAGR.[231] For the MIP, building observability into every agentic deployment is the technical prerequisite for every governance commitment the audit makes, and the commercial foundation of the monitoring relationship that makes the governance audit a recurring revenue engagement rather than a one-time project.

For channel partners serving clients in regulated industries or operating across jurisdictions: healthcare, legal, financial services and any business with Canadian, European or multinational data exposure; the governance audit must address a dimension that most current AI deployments have not yet confronted: where AI systems physically process and store data, and which legal frameworks govern that processing. Sovereign AI encompasses how intelligence is created, trained and deployed across infrastructure, models and applications; making it broader in scope than data residency alone, and spanning four distinct dimensions: territorial (where data and compute physically reside), operational (who manages and secures them), technological (who owns the underlying stack), and legal (which jurisdiction governs access and compliance).[232] Gartner predicts that by 2027, 35 percent of countries will be locked into region-specific AI platforms using proprietary contextual data, with nations potentially needing to invest up to one percent of GDP in AI sovereignty infrastructure by 2029.[233] For an SMB client whose AI agents are processing client health records, financial data or legal documents, the question of where that processing occurs and under which jurisdiction it is governed is a present-day compliance concern. The MIP who surfaces it, and who can advise on data residency configurations, regional compliance frameworks and sovereignty-aware deployment architecture is providing value that no horizontal AI platform and no generalist provider can replicate without the same domain depth.

THE REVENUE ARCHITECTURE OF THE AUDIT PRACTICE

The governance audit creates a revenue architecture that is structurally different from and more durable than initial AI deployment engagements. Deployment engagements are project-based: defined scope, timeline, deliverable, natural conclusion. Governance engagements are continuous: the AI environment the audit discovers and remediates on day one will have changed by day 90, as employees adopt new tools, vendors embed new AI capabilities in existing products and the client's own agentic deployments expand the surface that requires oversight.

The ScalePad 2026 MSP Trends Report documents the performance differential between providers who have built advisory practices around ongoing client relationships and those who remain primarily transactional. Among top-performing MSPs, 42 percent offer vCIO or strategic advisory services, compared to 29 percent of the overall MSP population.[234] The AI governance audit is the natural entry point for that advisory relationship; a structured engagement that begins with a concrete deliverable, transitions into an ongoing monitoring and optimization function and positions the provider as the strategic intelligence layer that keeps the client's AI environment compliant, secure and expanding in the right direction.

The technology providers who build this practice are building the client relationship that every other service line in their portfolio depends on. The SMB whose AI environment has been audited, governed and integrated into a managed intelligence framework by a trusted partner is not a candidate for competitive displacement by a horizontal AI platform or a born-agentic challenger. They are a client whose technology architecture is inseparable from the partner who built its governance layer, and that inseparability is the most defensible competitive position available in the managed intelligence market.

The agentic workforce economy does not wait for governance to catch up. The question now is which technology provider will have built that capability before their clients discover they needed it?

THE CONVERSATION THAT OPENS EVERY DOOR

The governance audit has one additional strategic property that distinguishes it from every other entry point in the digital labor portfolio: it begins with clients the provider already has, using information the provider can gather from existing monitoring and management relationships and it creates urgency that the client already understands because the risk is already present in their environment.

The deployment conversation requires the client to imagine a future state they have not yet experienced. The governance audit conversation requires only that the client acknowledge a present state they are already living in; one in which their employees are using AI tools the business has not reviewed, processing data the business cannot track and taking actions the business cannot audit. That conversation does not require a technology demonstration or a competitive displacement argument. It requires only the willingness to show a client what is already happening in their environment and to offer a credible path to managing it.

The providers who lead with that conversation in 2026 will capture new revenue lines, but they will also establish the governance architecture that makes every subsequent digital labor investment safer to make, faster to deploy and more defensible to maintain. And they will do so from a position of institutional knowledge that no competitor entering the relationship later can replicate.

The agentic workforce economy does not wait for governance to catch up. The question now is which technology provider will have built that capability before their clients discover they needed it?

229 PwC. "AI Observability for Enterprise AI Agents." PwC, 19 Feb. 2026.

230 PwC. "AI Observability for Enterprise AI Agents." PwC, 19 Feb. 2026.

231 Mordor Intelligence. "Agentic AI Monitoring, Analytics, and Observability Tools Market." Mordor Intelligence, 2025.

232 McKinsey & Company. "The Sovereign AI Agenda: Moving from Ambition to Reality." McKinsey, Dec. 2025.

233 Gartner. "Gartner Unveils Top Predictions for IT Organizations and Users in 2026 and Beyond." Gartner, 21 Oct. 2025.

234 ScalePad. "2026 MSP Trends Report." ScalePad, 2026.

Conclusion:

The Organizations That Define the Next Model

The macroeconomic signal is real and measurable.

The report has made a series of empirical arguments that bear restating as a unified picture before any strategic conclusions are drawn from them.

The macroeconomic signal is real and measurable. Labor productivity has risen at a pace not seen since the early years of the internet economy, driven by industries where AI-driven time savings are highest. The businesses capturing those gains disproportionately are the ones that have moved beyond tool adoption into full workflow integration; the 10 percent of SMBs that have completed the integration journey, generating the 111 percent profitability uplift that Deloitte's modeling identifies as the return available at the fully enabled stage. The 90 percent that have not completed that journey are not simply leaving value on the table. They are watching the competitive gap between themselves and the businesses that have widened each quarter, at a compounding rate that grows harder to close the longer the deferral continues.

The time savings are already accumulating and largely going un-reinvested. LOB workers using AI are recapturing 3.1 hours per day. IT workers are recapturing 3.6. Fewer than one in five SMBs has a deliberate strategy for where that recaptured time goes. The businesses that reinvest it into higher-order work (client relationships, product development, strategic planning) are building advantages their competitors cannot see yet. The businesses that absorb it passively are simply running the same operations at marginally lower cost, without the compounding organizational capability that deliberate reinvestment produces.

The risk dimension is not separable from the productivity dimension. Every AI agent deployed, automated workflow activated, data environment rationalized for AI consumption creates capability on one axis and exposure on another. The businesses that understand this are building more durable competitive positions, because the organizational discipline required to achieve deep AI integration is precisely the organizational discipline required to govern it safely. The businesses that do not understand this are accumulating a liability they cannot see — one that becomes visible only when a breach, a compliance failure or an ungoverned agent action makes the governance deficit suddenly and expensively apparent.

The channel is bifurcating. The technology partners who are building AI-native practices (deploying AI internally first, packaging their domain expertise into vertical playbooks, leading with governance rather than tools, positioning themselves as managed intelligence providers rather than managed service providers) are capturing premium pricing, longer engagement timelines and switching costs that compound with every deployment. The partners who are treating AI as an add-on to existing service models are watching margins compress as AI-native challengers redefine the value proposition of the category from below.

THE FIRM THAT IS BEING BUILT

This report documents an architectural transition; one that most of the organizations living through it have not yet named. The traditional SMB is organized around human routing. Information arrives through inboxes, meetings, tickets, spreadsheets and chat threads. Work moves from person to person. Context is carried manually. Approvals live in email. Employees function not only as decision-makers but as the transport layer that moves tasks, information and status across the business. The organizational chart is built around that transport function as much as it is built around the work itself.

The agentic SMB will be organized differently. Persistent agents will monitor workflows, maintain state, assemble context, coordinate systems, escalate exceptions and push decisions to humans only when judgment, authority or relationship is genuinely required. The employee-facing operating layer will increasingly be a an always-on agent that understands the employee's role, permissions, priorities, and patterns; actively managing open loops, coordinating downstream task-specific agents and prompting the human when needed rather than waiting to be prompted. The visible architecture may look like one employee working alongside one agent. Beneath that surface, that personal agent may be orchestrating dozens of specialized agentic functions across billing, collections, compliance, scheduling, client communication and internal knowledge work simultaneously.

The organizations that merely insert AI into existing workflows will capture the gains available in the first half of the J-curve. The organizations that redesign around this architecture will operate on a different curve entirely; one where human attention concentrates at the judgment, exception, approval and relationship points that genuinely require it, and where everything else runs.

The organizational architecture that emerges from this transition has three distinct layers, each with a different function and a different ownership model:

- **The agent layer** is where autonomous execution happens; the persistent personal agents, the task-specific workflow agents, the specialized back-office systems handling billing, collections, compliance, scheduling and client communication. This layer executes, monitors and coordinates. It operates continuously, without fatigue, and scales without headcount.
- **The human layer** sits above it, as the governing intelligence of the system. Humans own judgment calls, exception handling, relationship management, strategic decisions and the approval points where autonomous action requires human authority. In the agentic SMB, the human layer is smaller relative to output than in the traditional SMB, and substantially more productive per person, because human attention is concentrated where it is genuinely irreplaceable rather than distributed across work that agents can handle.
- **The governance and data layer** sits beneath both; the infrastructure that makes the other two layers function reliably. Clean, connected data environments. Access controls that define what agents can see, touch and act on. Policy frameworks that govern autonomous decision-making within defined boundaries. Observability systems that provide continuous visibility into what agents are doing and audit trails that prove it. Without this layer, the agent layer is ungoverned capability. Without the human layer above it, the governance layer is policy without judgment.

The three layers function as a system, and the managed intelligence provider is the architect and operator of all three.

That is the redesign this report has been describing from the first chapter. The productivity data, the valuation implications, the governance requirements, the vertical playbook strategy and the governance audit opportunity are all expressions of the same underlying transition: the small business economy is moving from human routing to agent routing, and the organizations that understand this as an architectural question rather than a technology question will define what comes next.

THE AGENTIC SMB OPERATING ARCHITECTURE

WHAT THE STRATEGY REQUIRES
The strategic argument of this report is not complicated. It has three components.

The first is sequence. The operations-first adoption pattern is not a limitation to be overcome. It is the correct order of operations. The springboard is in the back-office functions where digital infrastructure already exists, data is cleanest and outcomes are most measurable. The governance audit begins with clients already in the portfolio, addressing AI tools already running in their environments. The vertical playbook starts with the two or three industries where existing client concentration and accumulated domain knowledge are deepest. Every ambitious deployment begins with the foundational work that makes it reliable: data readiness, access governance, workflow definition, policy architecture. The providers who respect this sequence close deployments. The ones who skip it generate the pilot purgatory statistics that Omdia and McKinsey are finding across the channel.

The second is compounding. The managed intelligence practice should be considered a compounding asset. The second vertical playbook deployment is faster than the first because the domain knowledge from the first is encoded in the configuration, governance framework and outcome benchmarks. The tenth governance audit in a legal services practice is faster and more defensible than the first because the exception categories, compliance boundaries and data architecture patterns are understood deeply enough to have been systematized. The switching costs that accumulate across a portfolio of deeply integrated clients are not a lock-in strategy; they are the natural consequence of expertise that cannot be replicated from a standing start. Providers who build this understand that every engagement is an investment in the next one, and they price and sequence accordingly.

The productivity returns documented here are the early signal of a deeper architectural shift.

The third is urgency. The window in which the vertical playbook, governance audit and internal-first flywheel represent differentiated capabilities rather than table stakes is not indefinitely open. Gartner's projection that fewer than five percent of enterprise applications are currently integrated with task-specific AI agents, rising to 40 percent by end of 2026, is not a prediction, but a description of the deployment velocity that is already underway, in the clients that providers are managing today, with or without the governance architecture required to make that deployment safe. The providers who have built the audit practice, playbook library and governance frameworks before that wave arrives are the ones positioned to guide it. The ones who are still building those capabilities when clients begin demanding them are competing from behind in a market that is moving faster than they are.

The urgency argument has a sharper edge for technology partners reading the same data through the lens of their own client portfolios. Pax8's discovery research, interviewing SMB decision-makers across in-house, outsourcing-curious and already-tethered segments, found that in the segment most directly relevant to the channel, SMBs already working with a managed service provider, only one in four had an MSP relationship in which AI was being delivered through the partnership rather than around it.[235] The remaining three were either having the AI conversation independently or pushing the MSP to join the conversation already underway. The structural implication is consistent across the report. The window is not closing because AI capability is becoming commoditized. It is closing because the SMB is building the AI strategy that the MIP transition was meant to anchor. The providers who arrive late to that strategy will not be the leading advisor on the work that defines their client's next decade.

The agentic workforce economy does not wait for anyone to be ready for it. It is the current competitive environment. The organizations, SMBs and their technology partners alike, who treat it as such will define the next operating model for small business growth.

THE TWO AUDIENCES THIS REPORT IS WRITTEN FOR
The SMB that has read this report should take away a single operational conclusion: the productivity returns documented here are the early signal of a deeper architectural shift, and the businesses that treat this moment as an invitation to redesign how work flows are the ones who will arrive at the organizational model this report has described. The primary barrier between current state and that model is not technology.

The technology provider who has read this report should take away a different conclusion: the clients who need that organizational architecture most urgently are not future prospects, but current clients, running ungoverned AI tools in their environments right now. They are accumulating governance debt with every new employee adoption, every new SaaS platform with an embedded AI capability, every agentic deployment made without a policy framework to contain it. The revenue opportunity this report documents is sitting in the existing portfolio, visible to any provider willing to look for it and equipped to address what they find.

The agentic workforce economy does not wait for anyone to be ready for it. It is the current competitive environment. The organizations, SMBs and their technology partners alike, who treat it as such will define the next operating model for small business growth.

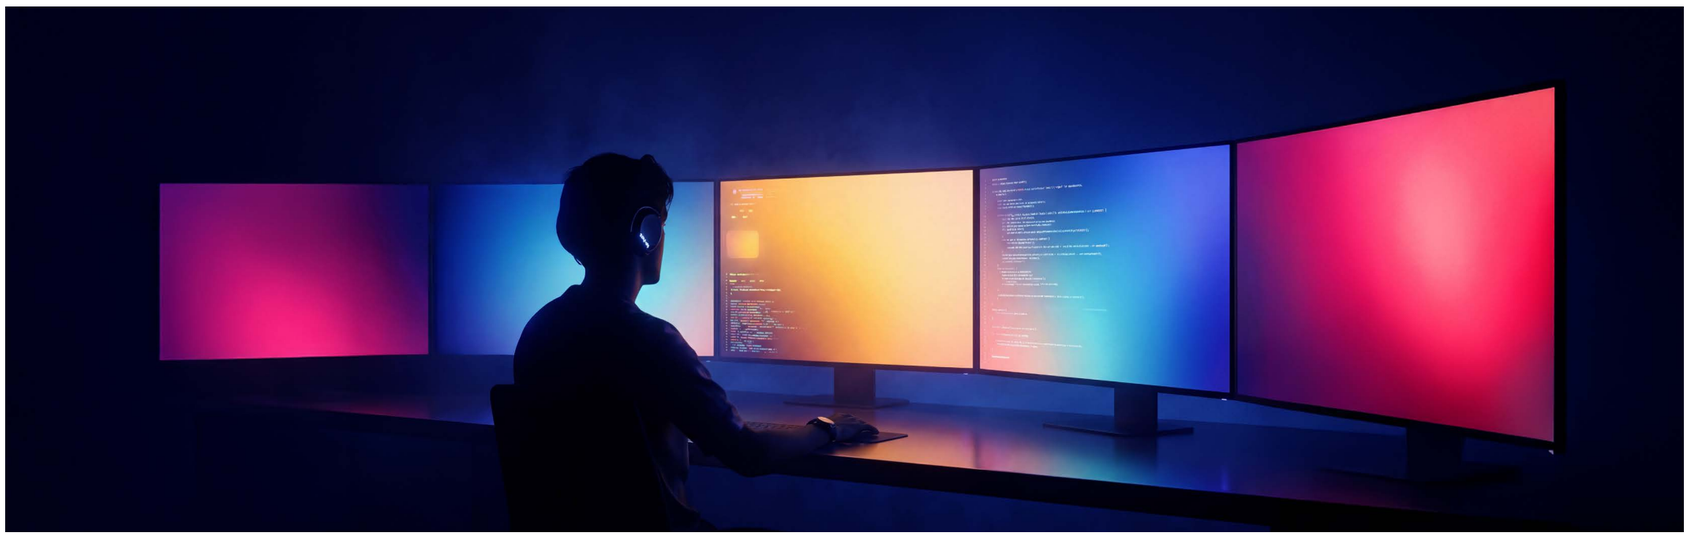

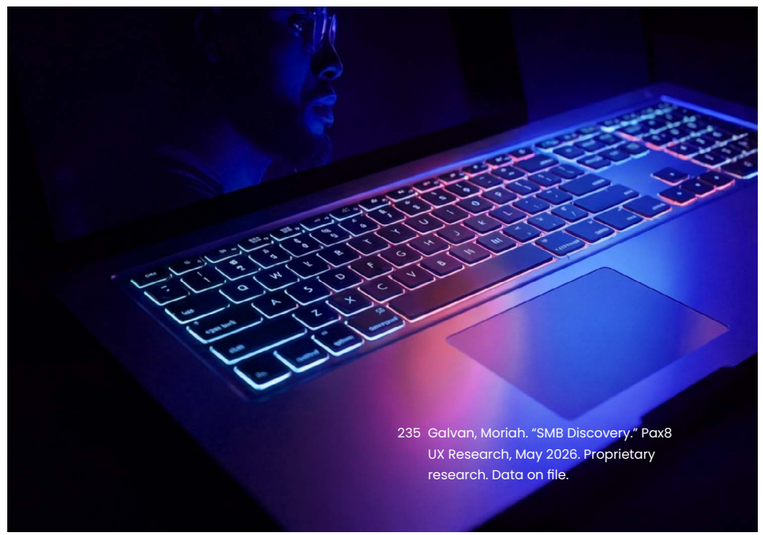

235 Galvan, Moriah. "SMB Discovery." Pax8 UX Research, May 2026. Proprietary research. Data on file.

The future of business will not be built by legacy systems or human-scale operations.

It will be built by autonomous intelligence, guided by the vision and trust of those who understand their customers best. In this new world, Managed Intelligence Providers are not optional participants but essential enablers of business transformation. The relationships built over decades, grounded in service, trust and proximity, will become their greatest assets in the years ahead. The next decade will not simply belong to those with the biggest budgets, but to those with the biggest ideas and the networks of partners and AI agents needed to execute them.

We are witnessing the convergence of trends that lower the barriers to building: it is easier to start a business, easier to access cutting-edge technology, and easier to scale rapidly with AI doing the heavy lifting. An SMB with a powerful idea can truly become a market maker overnight. Opportunity is no longer limited by size, but by imagination and execution. Those who move quickly and think differently will define what comes next. Momentum now favors the bold, the curious and the connected. The tools are already on the table; it is time to build.

The Agentic Workforce Economy explains the ***why.*** The MIP Playbook shows you the ***how.***

Schedule a Pax8 demo.

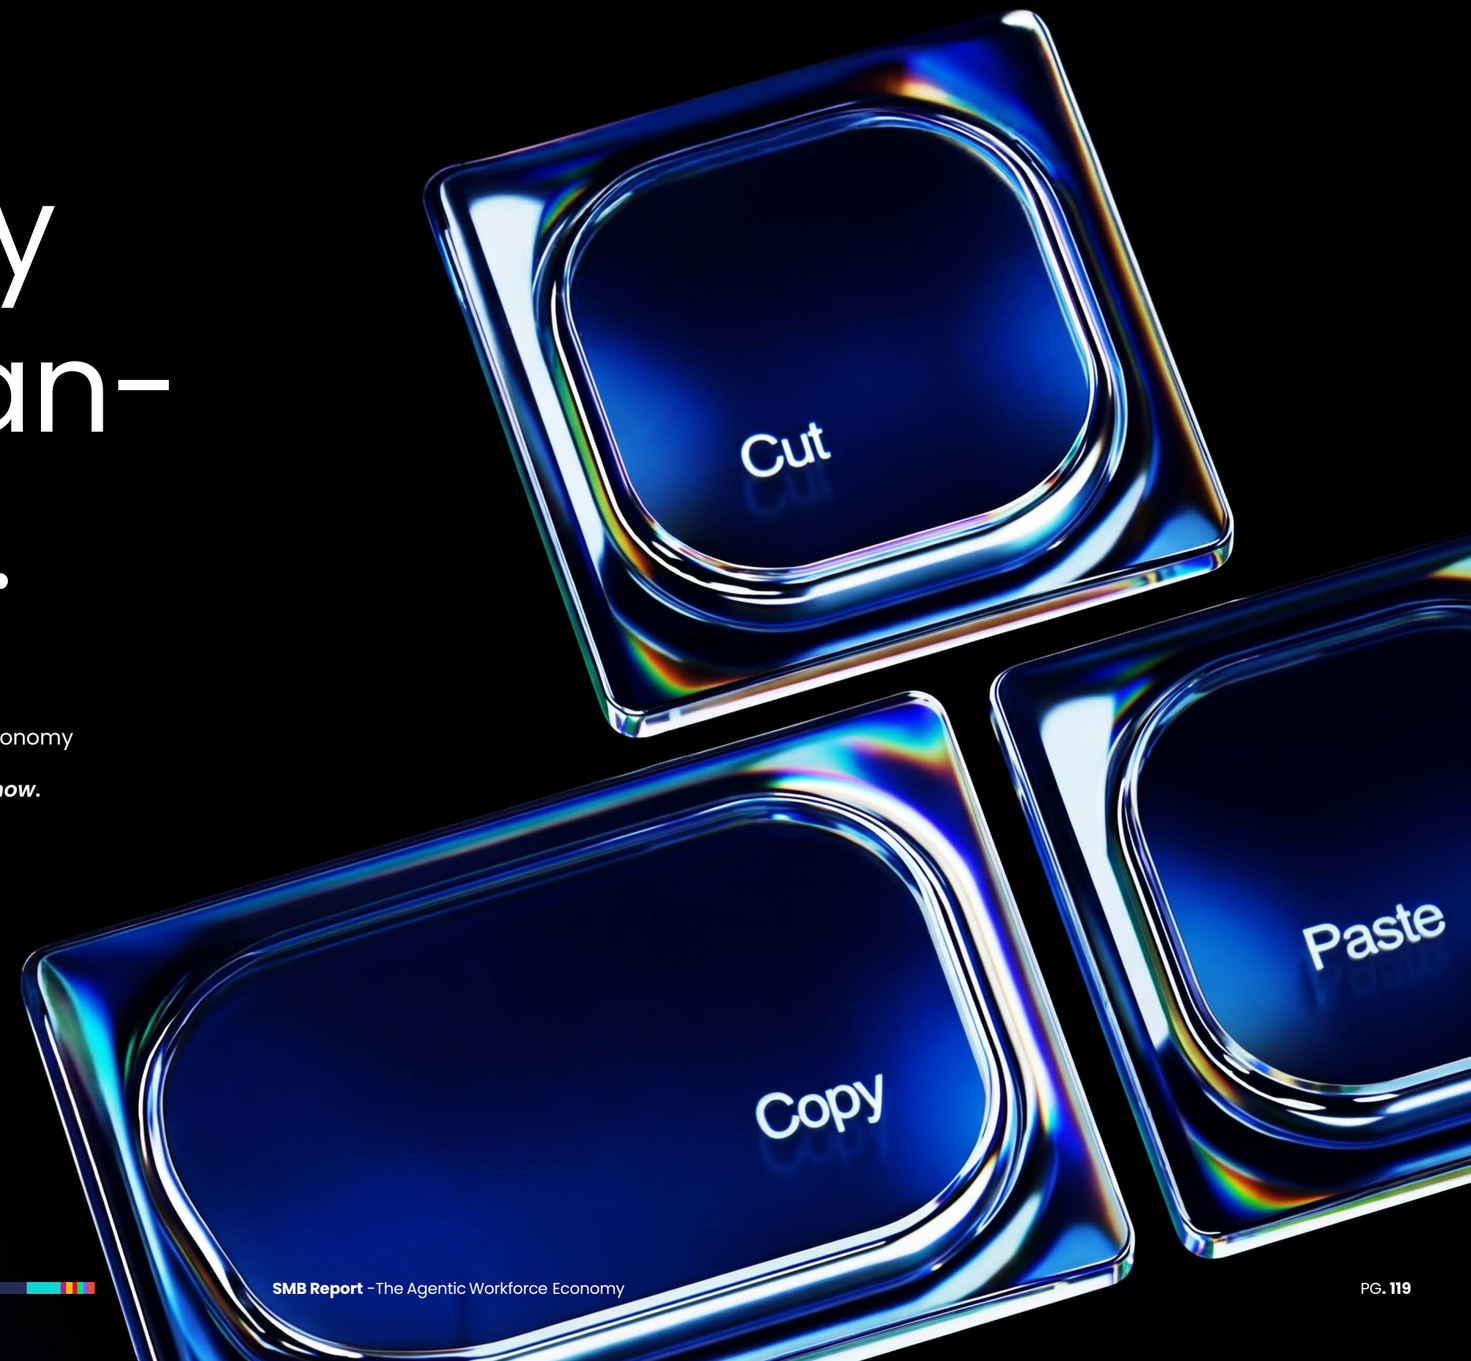

Thank you