

# Breaking the copilot ceiling

Unlocking ROI from genAI copilots to enable an agentic future

# Contents

- 1 Has the enterprise battle already been won?**
- 2 The tipping point of adoption**
- 3 The workforce to create value from copilots**
- 4 Concluding thoughts**
- 5 Reflections in practice**
- 6 References**

# Executive summary

Despite widespread adoption of Generative AI copilots—especially Microsoft 365 Copilot—most organizations struggle to realize meaningful ROI. Research shows up to 95% of GenAI projects fail to deliver value, revealing a gap between technological capability and business impact.

The paper identifies three core blockers: weak data governance and security, unclear value realization, and low AI literacy. These challenges are compounded by fragmented tool usage and reliance on personal AI subscriptions, which undermine enterprise adoption.

Kubrick introduces a solution: AI Business Partners—consultants trained to bridge the gap between technology and business processes. Embedded within teams, they develop secure, scalable copilot use cases with measurable ROI. Case studies across banking, private equity, and consumer goods show how targeted deployment can unlock value and accelerate adoption.

To break the “copilot ceiling,” organizations must shift from ad hoc experimentation to strategic implementation. With rising cloud costs and competitive vendor pricing, the window to build foundational value is closing. AI Business Partners offer a practical path to scale copilots effectively and prepare for the agentic AI future.

# Introduction

## Copilots & Agents - A Dichotomy of Hype

“We’ve entered the era of AI agents.” – Microsoft

“We envision a world in which agents operate across individual, organizational, team and end-to-end business contexts. This emerging vision of the internet is an open agentic web, where AI agents make decisions and perform tasks on behalf of users or organizations.”

This is the future, according to Microsoft’s Chief Communications Officer, Frank X. Shaw [1]. His vision is underpinned by the 230,000 organizations - including 90% of the Fortune 500 - who are beginning to build agents.

How close we are to this agentic future is the question at hand.

MIT’s Project NANDA [2] shares Microsoft’s mission, aiming to create a decentralized agentic web for the benefit of the economy and society at large. Their

# 95%

The State of AI In Business 2025 report made headlines in August, revealing that 95% of corporate Generative AI projects fail to demonstrate return on investment [3].

foundational research, however, has found a startling gap between the technological capabilities of Generative AI – the key building block of Agentic AI – and its impact.

Gartner’s 2025 AI Hype Cycle encapsulates this tension between the vision of an agentic future and the current reality of GenAI adoption: as AI Agents hit the Peak of Inflated Expectations, GenAI is descending into the Trough of Disillusionment (see fig i).

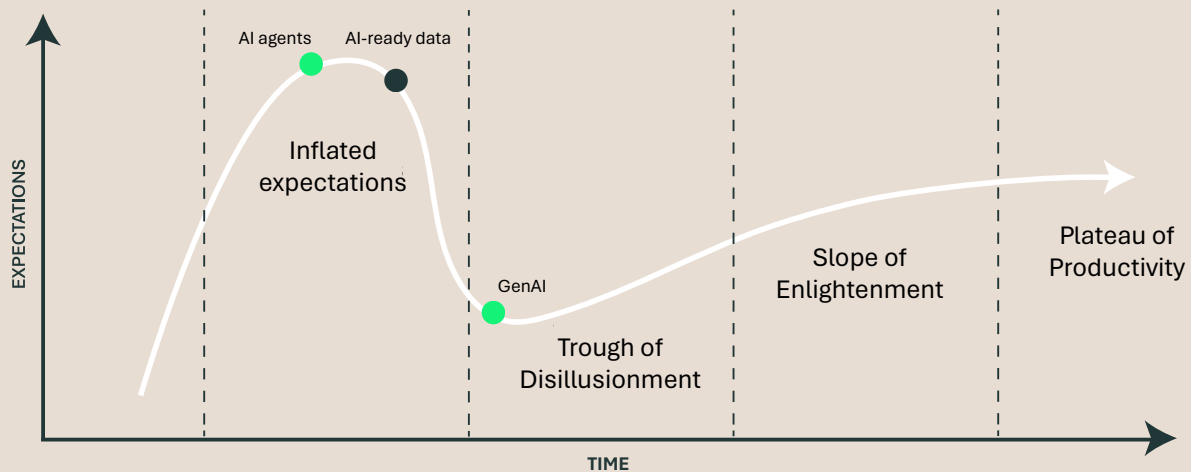


Figure i: Current position of GenAI and AI Agents on the Hype Cycle [Source: Gartner]

“With AI investment remaining strong this year, a sharper emphasis is being placed on using AI for operational scalability and real-time intelligence,” said Haritha Khandabattu, Senior Director Analyst at Gartner, commenting on the release of the Hype Cycle in June. “This has led to a gradual pivot from generative AI (GenAI) as a central focus, toward the foundational enablers that support sustainable AI delivery, such as AI-ready data and AI agents.” [4].

Despite difficulty proving value, the roll out of copilots continues to accelerate across all types of organizations,

from large enterprises looking to streamline at scale to small businesses seeking the competitive edge. In this new era of AI, it is still all to play for; tech giants, data platform specialists, and start-ups are all working overtime to cement their relevance in the market.

So, in this complex and competitive technology environment, what does it take to make the age of agents a reality? In this paper, we examine the current state of GenAI copilots to understand what organizations need to do to unlock their value - and pave the way for the agentic future.

# Has the enterprise battle already been won?

While ChatGPT was the disruptive force behind the rise of GenAI, Microsoft 365 Copilot is becoming ubiquitous in the workplace. This section examines the factors that influence the copilot marketplace to understand how we got here – and where we'll go next.

## ChatGPT vs M365: When legacy embraces disruption

When ChatGPT burst onto the scene in 2022, it gave rise to the universal understanding of a Generative AI ‘copilot’; as the fastest consumer app to reach 100 million users in history [5], it kickstarted the race to improve, accelerate, and integrate GenAI copilots into our personal and professional lives.

Three years on, OpenAI’s ChatGPT is still a dominating force in the GenAI marketplace, earning the colloquial collective noun and verb for using a GenAI tool: now, when scouring the internet for answers, one may ‘ChatGPT it’ in the same way that they may have once ‘Googled’ it in the age of search engines.

While ChatGPT is still prevalent as a personal productivity tool, in the workplace Microsoft’s M365 Copilot is coming out on top. In just the last year, M365 Copilot implementation in enterprises has surged past, with research suggesting a 11% lead ahead of ChatGPT corporate subscriptions (see fig ii).

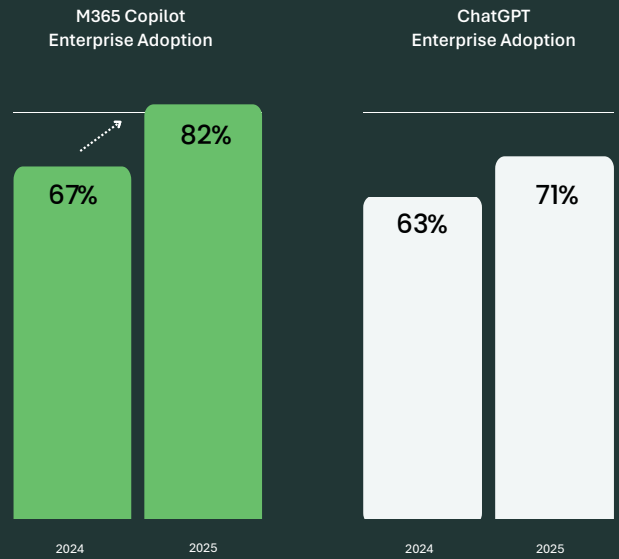


Figure ii: M365 Copilot is growing faster in enterprises than ChatGPT [Source: Jeffries via AlphaSense]

However, while Microsoft Copilot is set to become ubiquitous in the corporate world, there is appetite for variety in the market. Organizations are building layered, integrated GenAI ecosystems, as demonstrated by Figure iii - with M365 Copilot as the cornerstone.

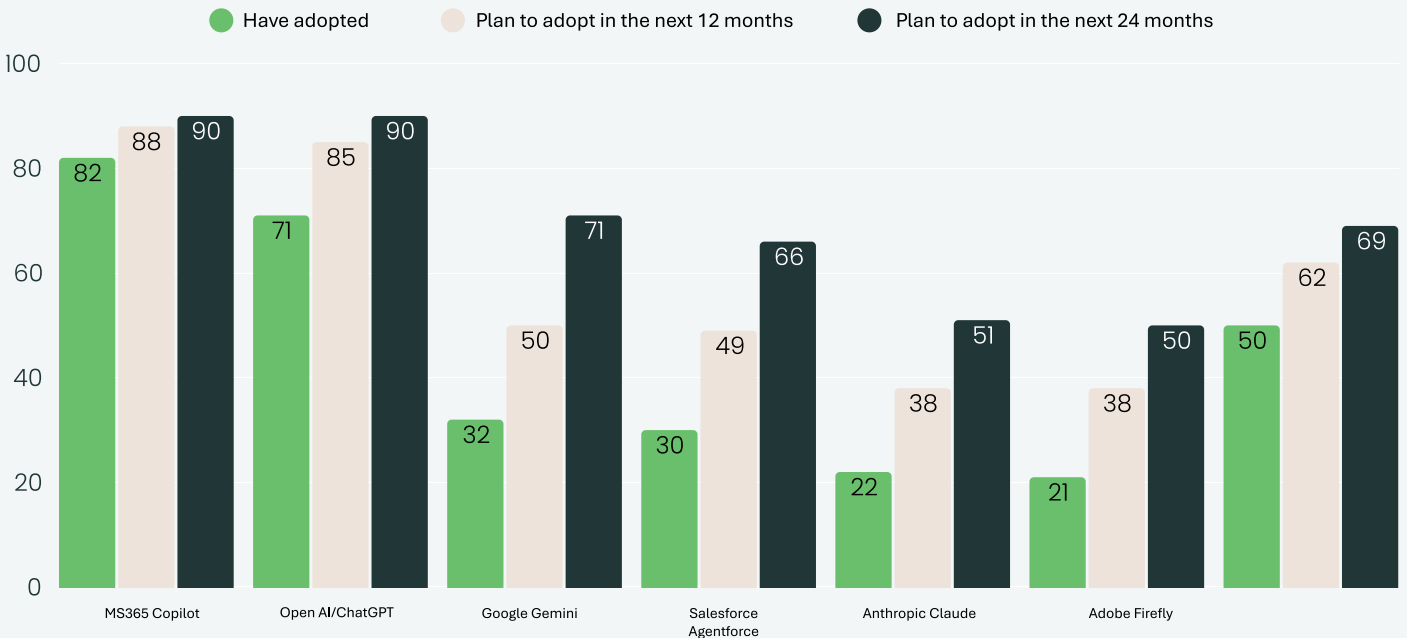



Figure iii: The adoption of GenAI tools continues to grow across enterprises [Source: Jeffries via AlphaSense]

## Why is M365 Copilot storming ahead?

The dominance of Microsoft in the GenAI copilot space is more about commercial benefit than technological. It can be attributed to three key factors:




### Integration with existing applications

Microsoft's stronghold over enterprise productivity software means instant integration with data from their everyday tools (i.e., Excel, Word, PowerPoint, SharePoint, and email). This gives users generated content and analytics that is tailored to their specific needs – with enhanced accuracy as a result.



### Pricing

At \$30 per seat, M365 Copilot is half the price of ChatGPT's Enterprise offering. Moreover, numerous business leaders have reported discounting opportunities in their enterprises to get them onboard [6], with major discounts like that for the US Government making headlines [7].



### Security

Staying in the Microsoft ecosystem provides an instant governance wrapper. Security concerns with inputting sensitive company information into external tools mean many organizations uphold outright bans of ChatGPT and others, despite the enhanced security and privacy of enterprise subscriptions which do not train the public, open-source model [8].

## Going to market with GenAI

This is not the first time Microsoft has used integration with their existing customer base to enter a new market – and win. When Microsoft launched Teams to compete against Slack in 2016 - and then played to the boom of Zoom during the COVID-19 pandemic - they soon claimed a large market share, thanks to the immediate integration as well as significant price and security incentives [9].

Going wide: in playing to both their wide market and wide array of integrated software, Microsoft are inherently positioning Copilot as a Jack-of-All-Trades kind of tool. It allows everyday users to quickly connect and generate information from across their personal data ecosystem but is generally accepted to be limited in areas of more technical challenge, such as graphic design, complex analytics, or code development.

As more products compete in the copilot arena, market share is not diminishing but instead widening as they refine their specialist features and positioning. For instance, Adobe is taking the lead in design generation with Firefly, while Anthropic's Claude is the firm favorite for code development [10]. The sprawl of tools in Figure iii represents how this exploratory phase of AI adoption manifests: individual users and teams find solutions with the specialist capabilities they require, creating an organic, ever-growing AI ecosystem.

Where Microsoft is playing is the wrapper: by enabling integration with leading enterprise software vendors like Adobe, Jira, Monday.com, Salesforce, and others, M365 Copilot becomes the 'one-stop-shop' to query data for insight, augment their workflows, and eventually deploy agents to unlock end-to-end process automation.

## The ROI challenge

Although copilots are positioned as a one-stop shop for AI assistance, the ever-growing market of differing capabilities and specialisms inherently creates an ROI challenge: a growing catalog of subscriptions and licences which drives up overall AI spend while making it difficult to track and attribute value from each.

This is compounded by the 'shadow AI economy': a phenomenon whereby continuing use of personal LLM accounts, such as ChatGPT and Claude, diminishes the adoption and attributable value from enterprise accounts like M365 Copilot, as reported by both Project NANDA [11] and Bloomberg. [12]

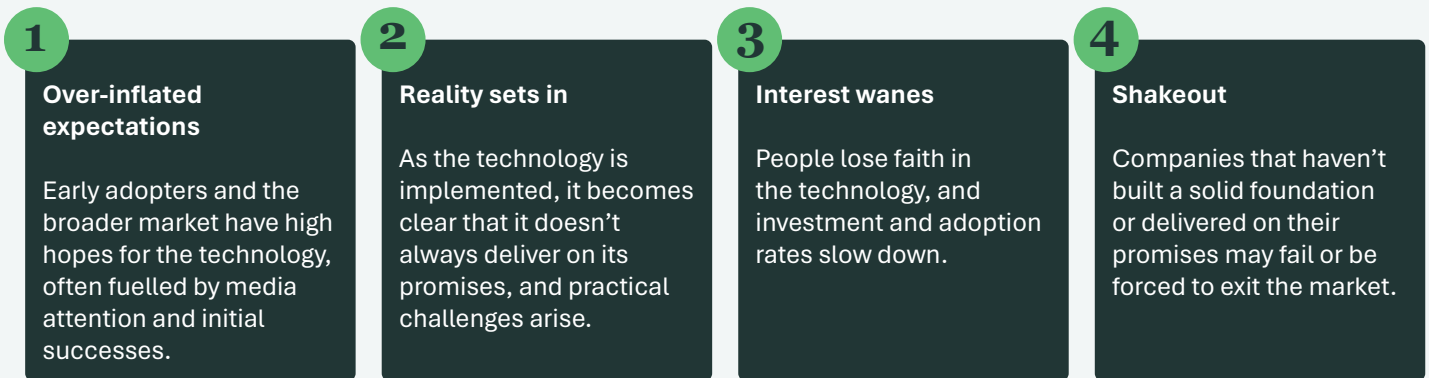
It is still all to play for in the GenAI copilot race. Microsoft Copilot may have market dominance across its enterprise accounts, but the measures of success are multi-dimensional. Undermined by tool sprawl and personal AI use, there is not yet a clearly defined conversion ratio of subscription to usage to value. This is the driving force that is sending Generative AI over the Peak of Inflated Expectations and into the Trough of Disillusionment.

# The tipping point of adoption

The peak of inflated expectations is behind us; the challenge of realizing value from GenAI still lies ahead. In this section, we unpack the challenges of putting copilots into production.

## Understanding the Trough of Disillusionment

As GenAI descends into the Trough of Disillusionment, we can examine how the phenomenon manifests to understand the root cause issues – and work to overcome them. Symptoms of the Trough of Disillusionment include:



Interest in GenAI is certainly not waning. Investment in tools is increasing as expectations for the agentic future remain high, as seen on the balance sheets of Big Tech: Google Cloud and Microsoft Azure revenues surged by 32% [13] and 39% [14] respectively in Q2 2025, fuelled by the demands of AI on cloud compute. Simultaneously, Meta credits its 22% revenue jump to the efficiency gains and margin strength made within their advertising - thanks to their own GenAI capabilities.

Market shakeout is minimal as vendors continue differentiating in capability and specialism to carve their niche in the market (see Part 1). Moreover, the biggest tech players are still vying with Microsoft for the generalist copilot market, leveraging pricing to create disruption. Google has now integrated Gemini features into all Google

Workspace Business and Enterprise plans, scrapping the \$20/user/month licence fee and instead nominally raising the overall Workspace subscription from \$12 to \$14 per month [15]. In doing so, they made a strong play for the Small Medium Business market while creating pricing pressure in the Enterprise space against Microsoft's \$30/user/month licence fee.

The predominant symptom of the Trough of Disillusionment for GenAI is the reality of implementation. Gartner's research into the blockers to deploying M365 Copilot found 3 main challenges that resonate across organizations: security and governance concerns, unclear ROI/value, and a lack of GenAI skills in workforce (see fig iv). These are the reality checks leaders need to assess their copilot readiness before they can move forwards.

### Leadership survey: the top blockers to deploying M365 Copilot

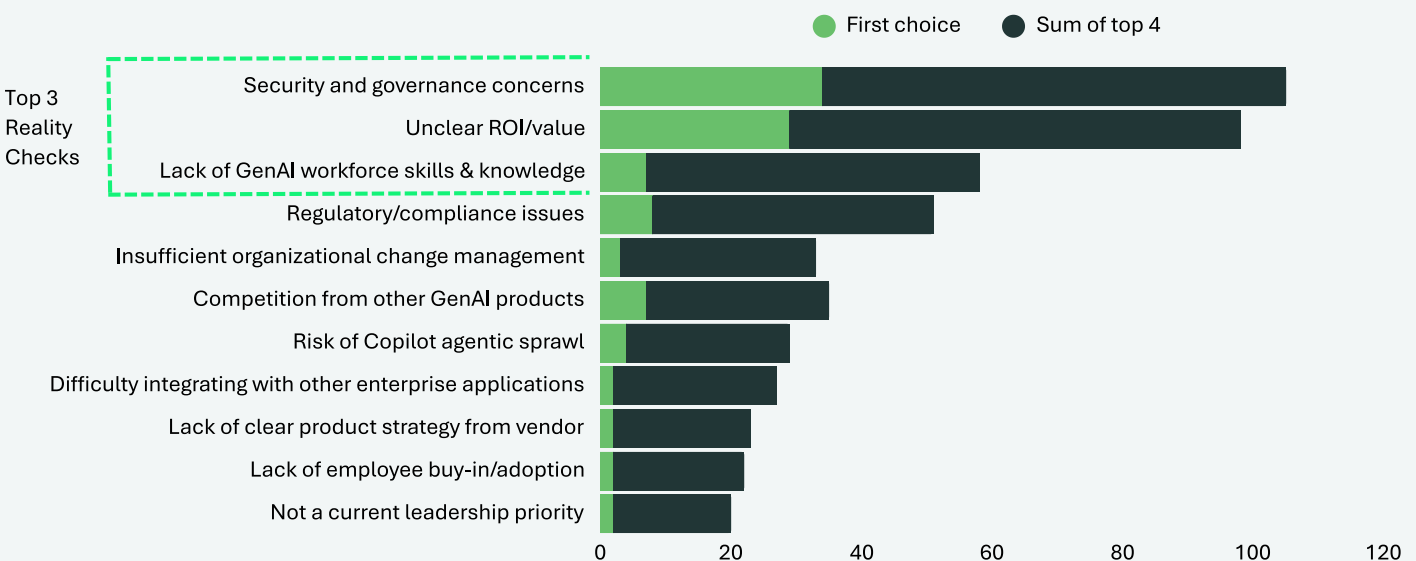


Figure iv: 2025 M365 Copilot Survey [Source: Gartner]

## Reality check 1: security and governance

In principle, M365 Copilot alleviates macro-level security and governance concerns about keeping sensitive data within the organization's ecosystem, rather than feeding a public model. However, there are still several security challenges blocking largescale rollouts of Microsoft Copilot [16]:

- Stale users: Nearly 88% of organizations have accounts that haven't been used in over 90 days, leaving sensitive data exposed to dormant identities.
- Over-permissioned cloud roles: The average environment contains thousands of overly permissive policies, making it difficult to enforce least privilege at scale.

- High-risk OAuth apps: One in four OAuth apps are high-risk and unverified, and over half of employees use them—often without IT oversight.
- Public links: 92% of organizations allow users to create public sharing links, which are frequently used to expose ungoverned and confidential information.

Moreover, most organizations do not have the data management practices to make trustworthy, accurate AI a reality. Gartner predicts 60% of AI projects in 2026 will be abandoned because of a lack of 'AI-Ready Data' [17], a term which they added to the 2025 Hype Cycle (fig ii). With cyberattacks on the rise [18], it is imperative that security and governance around any AI product, including pilots, are baked in from the start.

## Reality check 2: value realization

Research like Project NANDA has found that up to 95% of GenAI projects have failed to return any value [19] - and the outlook from the C-Suite is not much more optimistic. IBM's 2025 CEO Survey of 2,000 executives found only 25% of AI projects delivered the expected ROI [20]. The advent of copilot subscriptions has made the cost of AI more tangible: In a single enterprise organization, licencing fees scaled across thousands of employees quickly reaches the millions. Added to that the rising cloud costs of running these systems (which have been so

beneficial to Big Tech's revenue [21]), and you'll find CFOs examining a big dent against the bottom line.

The value in the top line is, conversely, less easy to tally up than its cost. Common practice is predominantly one-off or ad hoc tasks, like queries or basic analytics, which can be done in Copilot (or, more often in 'shadow' tools like ChatGPT). These personal productivity gains cannot be measured or scaled – and are fragmenting how teams work.

## Reality check 3: workforce requirements

For business users, the ability to work with copilots in natural language means it is quick to get started - in theory. The reality is that a general lack of GenAI literacy threatens to dichotomize employees into two types: 1) those who don't engage out of fear or reluctance and, 2) those who use without awareness of risks or adequate governance. Now that AI is being put directly into the hands of end-users, AI literacy is an increasing skill gap in the workforce that will make the difference between whether AI is successfully and safely adopted.

Moreover, the ease of plug-and-play copilots makes them inherently more difficult to prove value. The ability to input data of varying structures and quality, with little to no alignment against business processes results in moderate personal productivity - but no discernible impact on P&L [22]. In this way, GenAI copilots demand a new skillset in the workforce: data practitioners who can scope and develop use cases that return measurable, scalable value.

# The workforce to create value from copilots

**Copilots have the lowest technical barrier to entry than any AI technology – so why aren't everyday users seeing demonstrable value? The new greatest AI skillset is not technical skill, but the ability to understand its applications, commercial context, and ways to scale ROI fast.**


## Introducing AI business partners

Kubrick's founding principle is to create its own workforce skilled in the most pressing areas of data and AI – and the new age of copilots is no exception. Kubrick AI Business Partners are consultants who combine data engineering and governance with effective product delivery, specialized in copilot implementation, agentic workflows, and prompt engineering.

Embedded within teams, they help manage the entire product development and adoption lifecycle, including:

- Discovery of business pain points to frame Copilot use cases
- Development of lightweight business cases to support rapid development  
Connecting business processes with data structures through conceptual modelling
- Building and deploying intelligent agents using Copilot Studio
- Training and coaching business users in prompt writing and AI tooling
- Supporting change management by onboarding users and tracking adoption
- Working collaboratively with IT to ensure safe, scalable innovation

AI Business Partners are trained to overcome the 3 'reality checks' of copilot adoption with a guiding framework for their use case development: Working within business units, such as Finance, Marketing, or HR, they deliver high-impact, scalable use cases with:




### Improved security and governance

With contained systems, data sources, and monitoring, especially where data systems like CRM tools have more stringent data quality metrics than enterprise data.



### Clear value metrics and measurable ROI

Based on pre-existing workflows that use cases can augment and scale value across users, rather than random, one-off tasks or queries.



### Holistic AI literacy and upskilling

By focusing on a specific team and their requirements, rather than generic, hands-off training.

# Putting AI Business Partners into action

Realizing value from copilot goes beyond defining and executing strong use cases. It requires a critical understanding of technology ecosystems at large, strategic opportunities for scale, and the wider commercial objectives at play.

Here, we explore three macro-level business objectives which AI Business Partners help address while delivering value from copilot use cases.

## Objective 1: Creating an interconnected technology landscape with security

### Integrating M365 Copilot with third party tools

M365 Copilot is evolving a one-stop-shop by partnering with specialist tools like Adobe, Jira, and Salesforce (see Part 1). However, as copilots exacerbate the risks of over-permissioned access, stale users, and public sharing (see Part 2), integration with other tools must be designed with a security-first approach – especially where highly sensitive data is involved.

### Example: Microsoft + Adobe for Marketing

By integrating Copilot with the Adobe Customer Journey Analytics platform, Marketers can directly query their customers' digital engagement. AI Business Partners help design workflows with strong governance guardrails to extract data from Adobe CJA securely, so Marketers can automate campaign reporting and get recommendations for next actions, while the underlying customer data remains accessible to only authorised users.

## Objective 2: Scaling operational excellence

### Enabling use case replication across processes, teams, and functions

The quickest route to value from copilots is repeatability: the more users across business lines that can utilize a process flow, the greater the impact on margins – and adoption rates. AI Business Partners utilize frameworks to identify and prioritize the use cases that permeate across business functions, from Finance to Sales.

### Example: Copilot for Private Equity & Portfolio Companies

Much like individual business units within an enterprise organization, each Private Equity Portfolio Company operates in silo but will have many similar, applicable process flows which, if properly identified, scoped, and actioned, can realize a multiplying effect to the streamlining value of copilots. AI Business Partners centralize the knowledge and IP for PE firms and their Portfolio Companies to learn from each other and rapidly scale use cases to maximize return.

## Objective 3: Enhancing value from vendor relationships

### Making ROI from copilot a Win-Win

In the emerging copilot marketplace, aggressive pricing strategies (see Part 2) can dislodge vendor strongholds on strategic accounts – namely between Microsoft and Google. However, migrating to a new software suite is a drain of time and resource with a hefty price tag to execute. To counteract these pricing strategies, customers can look to their vendors and partners to help realize more value from their existing suite and prevent costly migrations.

### Example: Co-investing in the M365

Up against contract renewal deadlines, AI Business Partners can quickly scope and deliver use cases that prove the ROI from the existing productivity suite. As a Microsoft Partner, Kubrick can work in a co-investment model to fund AI Business Partners to prevent copilots becoming retired shelfware and drive longevity in existing partnerships while avoiding costly upheavals.

# Concluding thoughts

With a shared goal of creating an agentic future, the marketplace is evolving rapidly to enable promising agentic workflows. However, as individuals and organizations at large struggle to turn ad hoc use cases into measurable value, it is clear we have yet to break the copilot ceiling.

While Project NANDA gained notoriety for uncovering a discernible ROI on 95% of GenAI projects, what they instead illuminated was the value of the 5% that succeed. Their resounding messaging is simple: “land small, visible wins in narrow workflows, then expand” – a sentiment shared by all forms of AI tools and projects, from plug-and-play copilots to custom builds.

The technical risks of custom GenAI solutions are the strength of copilots: no lengthy product builds, no high-cost barriers to entry, no complex configurations that stall at Proof of Concept. Instead, the challenge of amassing value from copilots is a cultural one, testing organizational collectivism over individual interest. This kind of largescale cultural change must be underpinned by talent who can educate and evangelize, driving AI literacy with strong frameworks and process management across the enterprise.

In the infancy of the copilot marketplace, there are deals to be made as vendors battle it out for contracts and loyalty. As the major vendors enhance the top line with rising revenue from the cloud, Microsoft and Google’s pricing plays with discounting and integrations are giving organizations of all sizes a low-cost entryway into the technologies.

However, as the market matures and compute costs rise, analysts and business leaders warn of impending ‘sticker shock’ [23]. This moment provides leaders a window in which to build the foundation for value from copilots – and not get caught in a sunk cost fallacy. With the right skills, processes, and approach, businesses can scale copilots to provide the operational excellence that delivers exponential return, even as the market solidifies.

As copilots – and then agents – begin to augment work for individuals, teams, and organizations at large, there is a multiplying effect of scale. For the businesses which make headway now, the total of their streamlined processes will create an operational excellence and resilience that ushers in a new era of business for all.

# Reflection in practice

Breaking the copilot ceiling is more than a strategic imperative; it is a cultural and operational challenge that demands introspection.

Reflect on the ideas explored throughout this paper and apply them to your own context to help sharpen your goals, challenge points, and areas for improvement in your copilot strategy.

## Implementation & Ecosystem

- What scale is your copilot rollout? What examples do you have of successful and unsuccessful roll-outs?
- How does your organization's GenAI copilot ecosystem compare to the spread of tools in Figure iii?
- How does M365 Copilot or Google Gemini interact with other specialist tools?
- What are the technical challenges of integration in the ecosystem?

## ROI & strategy

- How are you identifying and prioritizing use cases? Do they scale across functions?
- Have you been able to identify and measure ROI in specific use cases?
- How are you measuring personal productivity vs process streamlining?

## Skills and approach

- How are you engaging and collaborating between technology teams and end users?
- What is your current workforce's GenAI skill level? What skills are missing when attempting implementation?
- What is your organization's GenAI literacy and cultural competency level? What skills are missing when attempting adoption?
- Do you have the necessary security and ethical protocols?
- Do you have examples of creating AI-ready data?

# References

[1] Shaw, FX (2025). Microsoft Build 2025: The age of AI agents and building the open agentic web  
Available via Microsoft

[2] Project NANDA (2025). The Internet of AI Agents  
Available via Project NANDA

[3] Challapally, A; Pease, C; Raskar, R; Chari, P (2025): The GenAI Divide  
Available via Project NANDA

[4] Gartner Newsroom (2025). Gartner Hype Cycle Identifies Top AI Innovations in 2025  
Available via Gartner

[5] Hu, K (2023). ChatGPT sets record for fastest-growing user base  
Available via Reuters

[6] AlphaSense (2025) Expert Insight: CEO Survey

[7] GSA (2025). Multi-Billion Dollar GSA OneGov Agreement with Microsoft Brings Steep Discounts for Microsoft 365, Copilot, and Azure Cloud Services  
Available via GSA

[8] Hays, S. (2025) ChatGPT Enterprise Connectors: How to Prepare for the Major Shift in Data Access  
Available via Varonis

[9] Curry, D (2025) Microsoft Teams Revenue and Usage Statistics (2025)  
Available via Business of Apps

[10] Kubrick Advisory Board (2025) Expert Insight: Chief Data Officer, Media & Communications

[11] Ibid 3

[12] Ghaffary, S. (2025) ChatGPT Is Wowing Office Workers, Posing a Problem for Microsoft  
Available via Bloomberg

[13] Alphabet (2025) Alphabet Announces Second Quarter 2025 Results  
Available via Alphabet

[14] Sophia, D; Nellis, S; Sonia, D (2025) Microsoft to spend record \$30 billion this quarter as AI investments pay off  
Available via Reuters

[15] Gownder, JP (2025) Google Pressures Microsoft 365 By Adding Gemini To Workspace For (Mostly) Free  
Available via Forrester

[16] Varois Systems (2025) Why least privilege matters more than ever  
Available via AlphaSense

[17] Edjlali, R (2025) Lack of AI-Ready Data Puts AI Projects at Risk  
Available via Gartner

[18] Greis, J; Sorel, M (2025) The cybersecurity provider's next opportunity: Making AI safer  
Available via McKinsey

[19] Ibid 3

[20] IBM Newsroom (2025) IBM Study: CEOs Double Down on AI While Navigating Enterprise Hurdles  
Available via IBM

[21] Ibid 13 & 14

[22] Ibid 3

[23] Medici, A (2025) AI tools are nearing their 'Uber' moment. It may mean sticker shock for businesses.  
Available via Biz Journals

