eBook

The AI Activation Playbook

Beyond deployment: Activating AI at scale

nexthink

01.
The AI paradox, bought, not used

Enterprises are investing heavily in AI, **yet many still struggle to turn that ambition into lasting, organization-wide impact.** According to The GenAI Divide: State of AI in Business (MIT & Fortune), 95% of enterprise AI pilot programs deliver zero measurable financial return, underscoring the disconnect between experimentation and execution. The gap between "buying AI" and "using AI" has become one of the defining challenges of modern digital transformation; where technology moves faster than culture, and real impact depends on guiding people, not just deploying tools.

Why this matters now

For most organizations, the race to scale AI is well underway, but few are prepared for what comes after deployment. Many are discovering that implementation alone doesn't deliver transformation; deploying AI is far easier than embedding it into everyday work. Gartner projects that by 2026, more than 80% of enterprises will have deployed generative AI APIs or applications in production environments.

But deployment alone doesn't equal AI success. A study from Gartner found that on average only **48% of AI projects ever make it into production**, and it typically takes eight months to move from prototype to production.

Moreover, as organizations scale, they are realizing that governance, safety, and user confidence matter as much as model performance. Gartner also claims that over 40% of agentic AI projects are projected to be canceled by 2027 due to cost, unclear value, or governance complexity.

Contents

Key signals of the adoption gap

- Many AI deployments remain pilot-scale and never scale enterprise-wide due to security concerns and risks
- IT and business units struggle to measure usage and impact, leading to a “black box” on ROI
- Users avoid or abandon AI tools because they don’t feel confident, safe, or skilled
- Some use AI through alternate, ungoverned paths, creating shadow AI risk

Why people, not platforms, decide the success of AI

Technology often gets all the attention in AI discussions: models, infrastructure, data, APIs. But real transformation depends on people. Even a world-class AI model won’t deliver value if few people use it or use it incorrectly.

This means enterprises need to move beyond “deploy AI and hope it sticks” models. They must engineer adoption and governance from day one. Simply giving users access isn’t sufficient.

What the paradox implies for strategy

- **Visibility is critical:** IT and leadership need clarity on which AI tools employees use, who is using them, how often, and to what ability
- **Risk control matters:** shadow use, data leakage, and improper prompts can undermine AI tool confidence and employee usage quickly
- **Guided adoption is essential:** training, nudges, walkthroughs, and coaching help users cross the activation gap with real-time best practices and guardrails
- **Sustainability and efficiency must be built-in:** unnecessary prompting and misaligned usage impose real computational and environmental costs (we’ll dive deeper later)
- **Amplify employee feedback:** capture employee sentiment and survey responses to ensure your efforts align with their needs
- **Closed-loop measurement is nonnegotiable:** correlation between guidance, usage, and outcome is the only way to prove impact

In sum, the AI paradox is clear: **organizations are buying AI, but many are failing to make AI part of everyday work.** The disconnect is not a failure of models, it’s a failure of adoption. To close that gap, visibility, guided action, and measurement must become first-class capabilities in any AI strategy.

> Those who can measure and steer AI will not just withstand the shift; they will define the next era of transformation. But that is only possible if businesses are able to track how every AI tool is being used across the organization, empowering employees to use them with confidence and precision, while giving leaders visibility into outcomes and productivity at scale.”

Pedro Bados
Chief Executive Officer, Nexthink

02.
Why AI adoption stalls

To understand why many AI initiatives never scale, it helps to dig into the human, organizational, and technical barriers. Below are the primary reasons adoption stalls, each of which can be addressed today:

1. Fear, uncertainty, and compliance anxiety

Users often worry: Is it okay to give the AI this data? Will I violate policies? Will leadership object? In highly regulated settings, those fears are amplified. This is especially true in organizations where AI governance is immature. The risk of inadvertent data exposure, IP violations, or noncompliance with internal or external regulations deters participation. Until employees feel confident the tool is safe to use, adoption will lag. Without clear guardrails and enforcement, many employees take a better safe than sorry approach and avoid the tool entirely.

2. Lack of skills or confidence

Even when users are willing, many don't know how to use AI, how to frame prompts, how to integrate AI into daily workflows, or how to validate results. This is particularly true for employees in non-technical or domain roles. Without hands-on coaching, experimentation, or just-in-time help, early frustrations (e.g. poor outputs or misaligned results) lead to abandonment.

3. Traditional training can't support

Classroom sessions, webinars, and static user manuals can't keep up with how fast AI evolves, or how differently each employee learns. By the time a training deck is finalized, the AI interface or workflow has already changed. Employees are left to experiment alone, often making mistakes or abandoning the tool altogether. True AI adoption requires in-the-moment learning and reinforcement that meets users where they work, adapting dynamically as the technology does.

4. Unaware of AI's practical value

Some employees simply don't see AI as relevant or beneficial to their work. If the value isn't clear, many treat the tool as optional or "nice to have," not mission-critical. Drawing an analogy to Grammarly, a popular AI writing assistant that helps users correct grammar, refine tone, and improve clarity, many experienced writers see it as redundant or robotic; after all, they believe they already "know" English. In the same way, professionals may reject AI because they believe their domain knowledge suffices. Often, they undervalue the AI augmentation, speed, insight, and creativity boost that can arise even for experts.

5. Poor experience or friction

If the AI tool is hard to access (e.g. multiple logins, disconnected from core apps), adoption falters. If integrating AI into users' existing workflows is clunky, they won't bother. Friction kills momentum. Compatibility with business systems, single sign-on, embedded UI, and seamless connectivity matter as much as model quality. AI tools not aligned to employees' needs.

Companies will block AI tools or not provide licenses out of an overabundance of caution. Organizations need to understand what employees want and why to best support them.

6. Strict AI committees block access out of fear of the unknown

Many organizations form well-intentioned AI governance committees to manage risk, but overly cautious policies can backfire. By restricting access or delaying approvals, these groups unintentionally discourage exploration and slow innovation. Employees turn to unsanctioned shadow AI tools to get their work done, creating more risk, not less. Real progress comes when governance shifts from control to enablement, providing safe frameworks that encourage responsible experimentation and build confidence in AI use across the business.

7. Redundant or wasteful prompting is a sustainability cost

Every inefficient or repeated AI prompt carries a cost, not just in compute time and latency, but in environmental impact. Frontiers suggests that reasoning-enabled models can generate up to **50 times more CO_2 emissions** than concise-response models, underscoring how the complexity of a single query can dramatically alter its carbon footprint. Google Cloud adds that the median **Gemini prompt consumes roughly 0.24 watt-hours of energy and emits about 0.03 grams of CO_2**, a seemingly small amount that becomes significant when scaled across millions of daily interactions.

When employees repeatedly “brute-force” prompt, tweaking wording again and again until they get the right output, these inefficiencies multiply. The result is wasted compute cycles, unnecessary emissions, and avoidable cost. Without proper guidance or coaching, users default to trial and error, turning what should be intelligent assistance into an invisible source of environmental and financial drag.

8. Governance and measurement gaps

Even if adoption is successful in pockets, many organizations lack ways to measure, govern, or iterate. A lack of visibility into which departments are thriving or struggling, which prompts produce the most value, or how to correct usage drift often leads leadership to conclude AI isn’t working and freeze further investment.

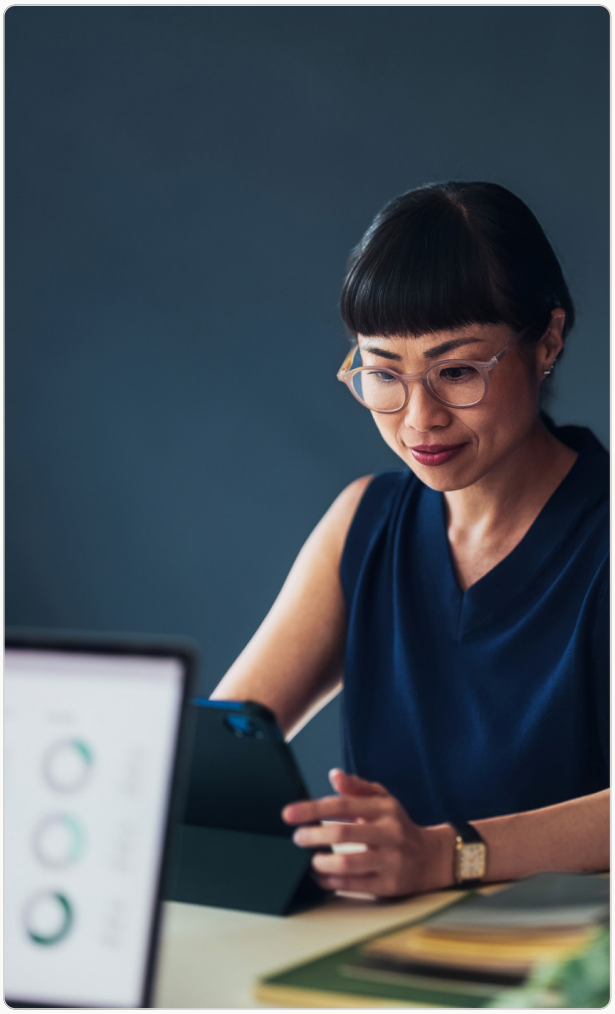

03.
The AI Activation Playbook powered by Nexthink

Most organizations understand that **AI adoption stalls**, but very few understand how to fix it. Despite enormous investment, most enterprises still lack a reliable, repeatable strategy for turning AI ambition into widespread, confident, everyday use. They deploy tools, run pilot programs, deliver initial training and then watch adoption plateau.

Without a clear operating model, AI success gets left to chance, enthusiasm, or isolated champions rather than sustained, measurable behavior change.

This is why Nexthink created the AI Activation Playbook, a practical, battle-tested framework that any organization can use to make AI adoption predictable instead of hopeful. It defines the steps, signals, guardrails, and feedback loops required to embed AI into real work, across every role and region, with confidence and compliance built in from the start.

Nexthink's AI Activation approach is not another training checklist or governance document. It is an **operating model for AI at scale:** a structured, human-centered system that unites visibility, enablement, governance, and continuous improvement into one motion. Rather than deploying AI and hoping employees adopt it, AI Activation engineers adoption deliberately, using real-time data and in-product guidance to drive confident, compliant use across the enterprise.

Nexthink powers this model end-to-end.

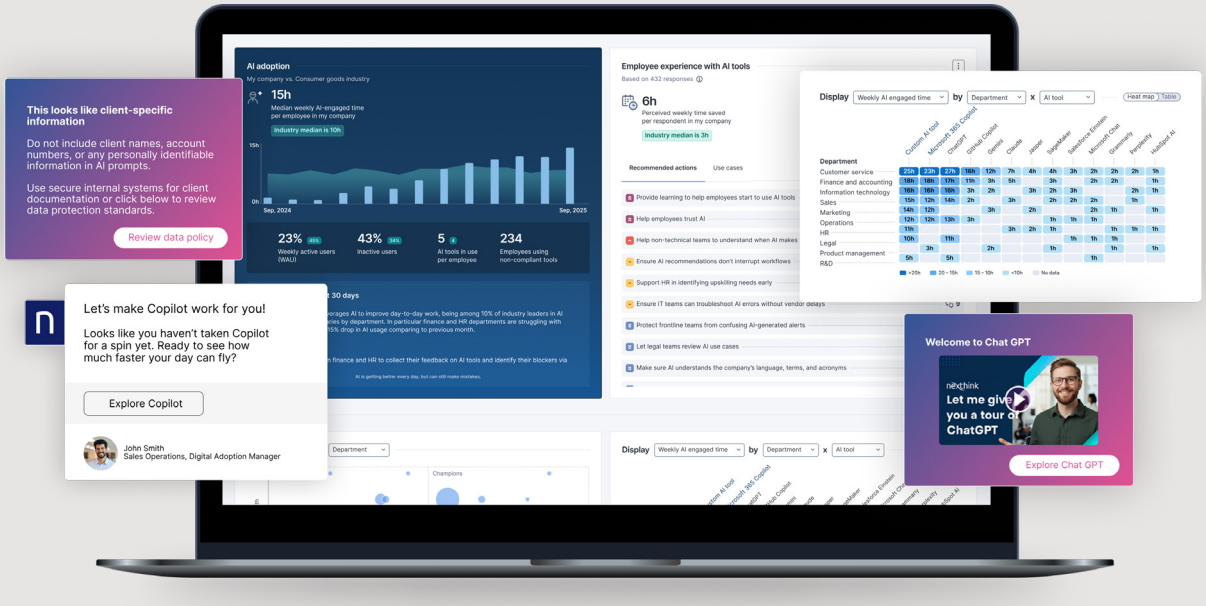

What is AI Activation?

AI Activation is the discipline of transforming AI tools from available to habitually used. It is the organizational muscle that ensures AI moves from experimentation to embedded practice.

At its core, AI Activation rests on five pillars:

1. Visibility into real behavior

You cannot improve or govern what you cannot see. AI Activation begins with telemetry on who is using which AI tools, how often, for how long, and with what perceived outcomes. Shadow AI must be surfaced, under-utilizers must be identified, and champions must be recognized.

2. Guided enablement in the flow of work

Employees need help as they work, not as they remember a training deck. Walkthroughs, nudges, and contextual coaching reduce uncertainty, build skills, and accelerate early success.

3. Risk-aware guardrails

Governance must be built into the workflow: upload warnings before data leakage occurs, prompt-level misuse detection before risky prompts are submitted, and redirect experiences that steer users away from unapproved AI tools.

4. Continuous feedback from employees

Adoption is shaped by confidence, perceived value, and real frustrations. AI Activation listens to employees through structured sentiment campaigns, allowing organizations to align interventions to their needs.

5. Closed-loop improvement

Every action must be measured. Did guidance increase active users? Did redirect experiences reduce shadow AI? Did fixing workflow friction improve adoption? Closed-loop analytics make AI sustainable, measurable, and defensible.

Nexthink is the **only platform** that unifies all five pillars into one end-to-end system.

How the AI Drive dashboard helps you gain visibility and build a strategy

AI Drive is the visibility engine behind AI Activation. It gives leaders what most organizations desperately lack: a single, trusted view of AI behavior across the enterprise.

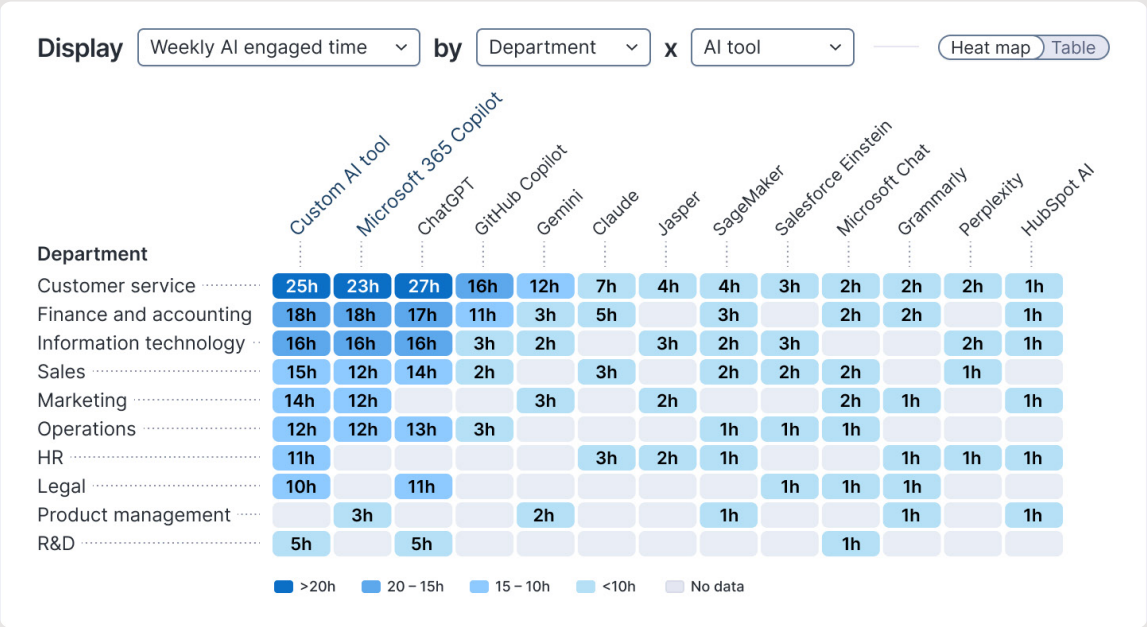

1. See AI usage across every tool, department, and persona

AI Drive captures real employee behavior across web-based AI tools, desktop applications, and Microsoft Copilot through deep integration with the Graph API. This gives leaders a multidimensional picture of how AI is actually used inside the business. Patterns emerge instantly: which tools dominate daily workflows, how much time employees spend engaging with them, and where adoption surges or stalls across departments. Champions, explorers, and under-utilizers become visible as distinct cohorts rather than vague assumptions, allowing leaders to understand the true distribution of AI skills and confidence across the organization. By illuminating what was once a complete black box, AI Drive replaces speculation with evidence and gives AI, IT, and transformation teams the clarity they need to make informed, strategic decisions about where tofocus enablement, governance, and investment.

2. Detect and quantify shadow AI

Shadow AI is often larger than sanctioned usage, but it becomes fully visible the moment AI Drive maps every AI tool employees are using across the estate. Once an organization designates its approved AI applications, everything else surfaced in the dashboard becomes, by definition, unapproved—and the goal is to drive that usage to zero. AI Drive pinpoints exactly where employees are veering into unsanctioned tools, while Nexthink Adopt can be used to redirect users to the right application: showing a pop-up when someone opens an unapproved AI app and steering them directly to the sanctioned alternative. This turns shadow AI from an invisible risk into a manageable, measurable behavior that can be read.

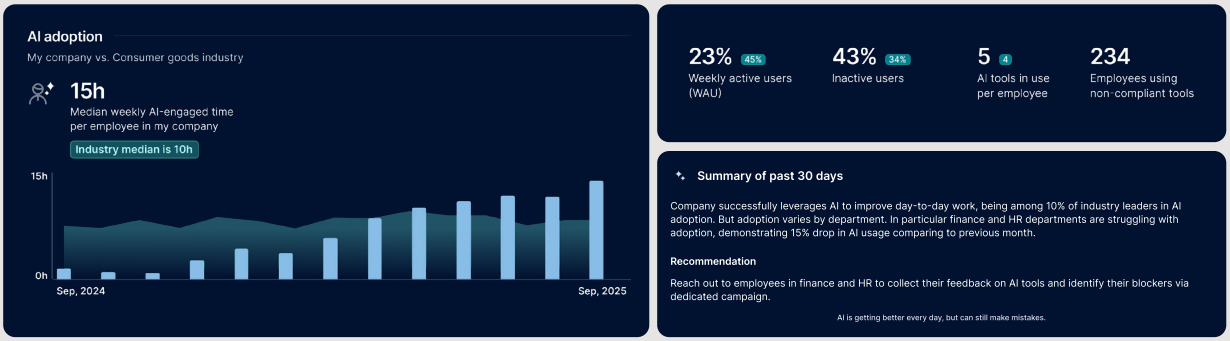

3. Benchmark performance

AI Drive's anonymized, cross-industry benchmarks place an organization's AI maturity in meaningful context. Instead of guessing whether adoption is "good enough," leaders can see exactly how their usage compares to similar companies, which tools demonstrate strong engagement relative to the market, and where adoption is lagging behind peers. This comparative lens sharpens strategy: it exposes underperformance early, highlights standout areas of success, and helps teams set realistic, evidence-based targets. Benchmarks transform AI adoption from an isolated internal effort into a measurable competitive position, making the case for focused investment far more compelling.

4. Understand employee experience with AI

True AI adoption cannot be measured through usage alone. AI Drive pairs behavioral telemetry with employee sentiment—capturing perceptions of time saved, frustrations with workflows, blockers, and overall confidence. This combination replaces anecdotal feedback with quantifiable evidence about how employees feel as they engage with AI tools. Leaders gain clarity on where people struggle, which workflows create friction, and where targeted guidance or redesign is needed to increase confidence and value. By unifying experience and behavior, AI Drive ensures AI programs evolve around real human needs rather than assumptions.

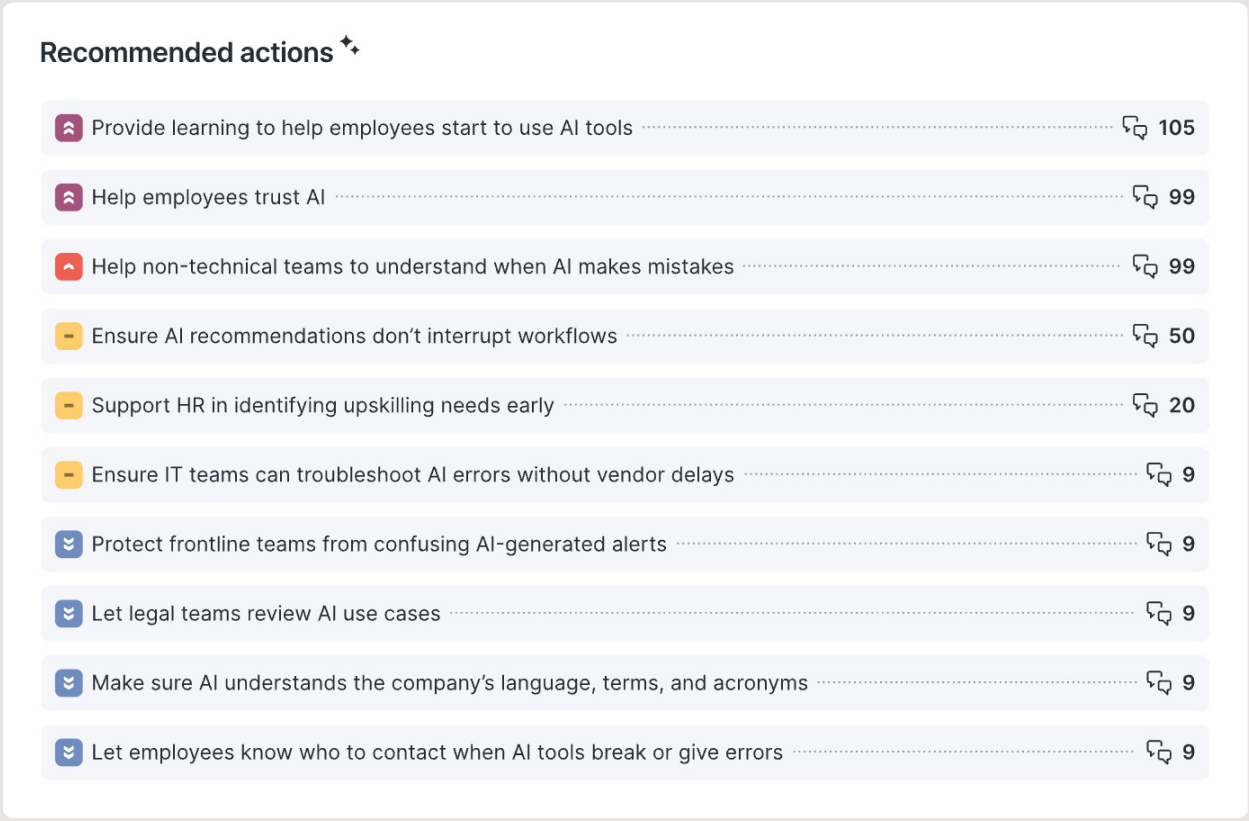

5. Prioritize actions with recommended next steps

With both usage and sentiment insights in hand, AI Drive automatically synthesizes thousands of data points into a prioritized set of recommended actions. These insights translate directly into operational improvements: enabling first-time Copilot users with guided walkthroughs, reinforcing policy guardrails around risky uploads, resolving performance issues that depress adoption, or targeting underutilizing cohorts with tailored enablement. For the first time, organizations receive a strategy that effectively writes itself—data-driven, sequenced, and tied explicitly to business outcomes. Instead of overwhelming teams with information, AI Drive delivers a focused plan for scaling AI safely, confidently, and efficiently.

From insight to execution: implementing AI Activation with Nexthink

Visibility gives organizations the insight to build an AI strategy, but strategy only becomes real when employees change how they work. This is where Nexthink's activation model comes to life. By combining in-flow guidance, targeted nudges, real-time guardrails, sentiment-driven refinement, and automated access, Nexthink turns insight into meaningful action across every department. The following scenarios illustrate how AI Drive and Adopt work together to steer behavior, build confidence, eliminate risk, and embed AI into daily workflows—creating a measurable lift in adoption and a more empowered, AI-fluent workforce.

1. Redirecting to the right AI tool for safe, unified use

Organizations often struggle with shadow AI because employees instinctively turn to familiar public tools like ChatGPT rather than the approved enterprise option. Nexthink solves this by detecting when a user opens an unapproved AI tool and immediately delivering a smart in-app pop-up that explains the risk and redirects them to Copilot. Employees receive instant guidance on why Copilot is the trusted environment, how to access it, and how to use it securely. As this behavior is reinforced, AI Drive shows a clear shift: shadow AI usage declines, sanctioned tools rise, and AI adoption becomes safer, more consistent, and fully measurable across the workforce.

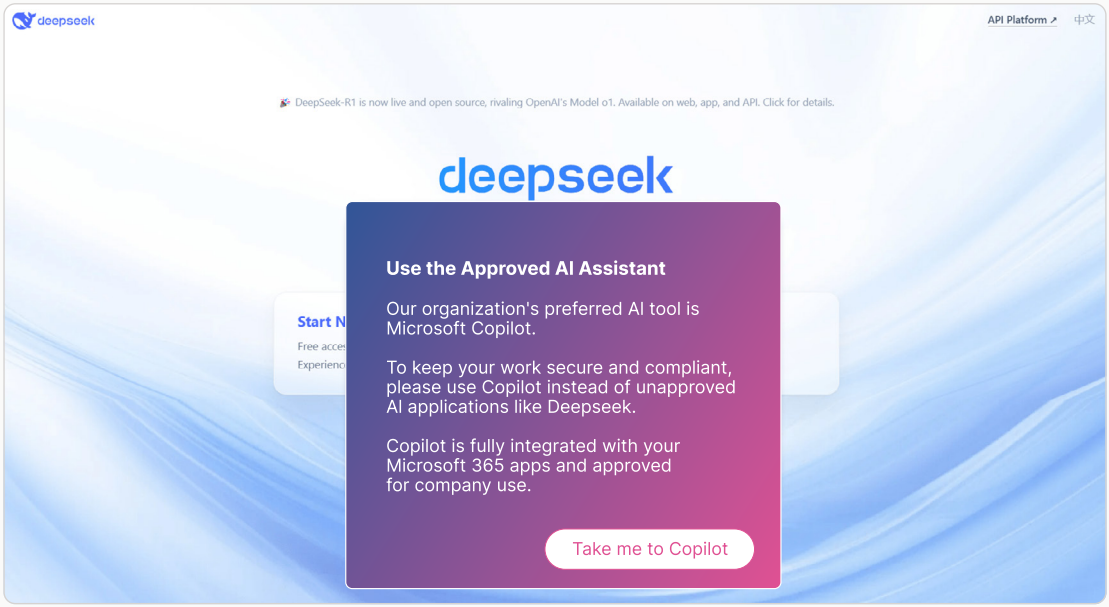

2. Guiding every generation of workers

AI confidence varies dramatically across a modern workforce. Some employees dive in eagerly, while others feel overwhelmed or uncertain. Nexthink bridges this gap through tailored walkthroughs that meet employees at their individual skill level. From AI-natives to late adopters, users receive step-by-step coaching embedded directly into their existing workflows, making Copilot feel intuitive rather than intimidating. Instead of generic training sessions, guidance appears exactly when people need support, helping them build trust and fluency quickly. This ensures equitable adoption across age groups, roles, and digital comfort levels, turning AI into a natural extension of everyday work for everyone.

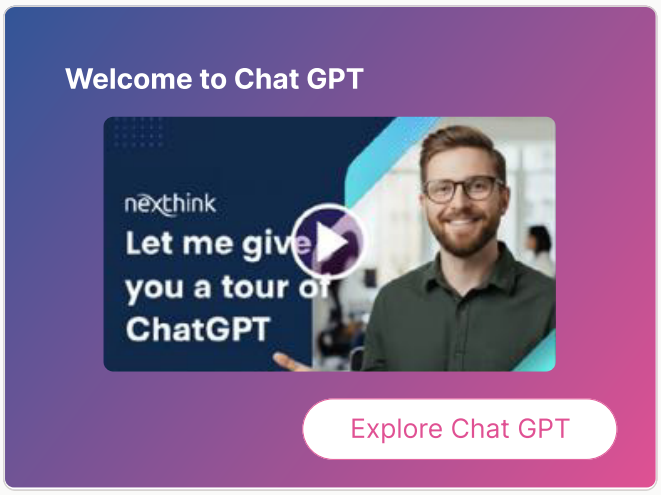

3. Driving confident AI adoption

Many employees are told to “use Copilot daily,” but without clarity or hands-on direction, they hesitate and adoption stalls. Nexthink transforms this experience by delivering clear, simple, in-context guidance that shows employees exactly how Copilot can accelerate their work. Nudges prompt users who have never engaged before, and real walkthroughs demonstrate practical workflows, instantly boosting confidence and familiarity. As employees begin to realize how much faster tasks become, engagement grows organically. Nexthink helps users move from avoidance to active usage, enabling organizations to close the adoption gap and ensure their AI investments translate into meaningful day-to-day impact.

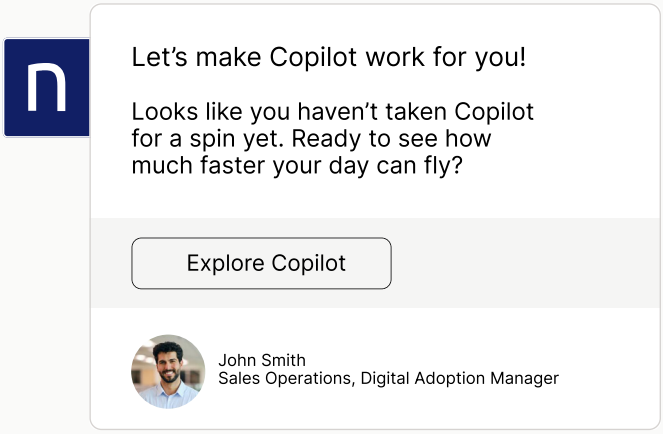

4. Training-gated access for responsible AI use

AI adoption only succeeds when compliance and confidence develop together. Nexthink enables this through training-gated access: new users encounter a welcome screen that requires them to complete mandatory learning before they can begin using Copilot. Once training is finished, Copilot unlocks automatically and a guided first task launches to build early success. This guarantees that every employee enters the AI environment fully informed about data handling, responsible use, and organizational expectations. The result is an ecosystemwhere leaders maintain compliance from day one, and employees start using AI with clarity, confidence, and the right foundational skills to succeed.

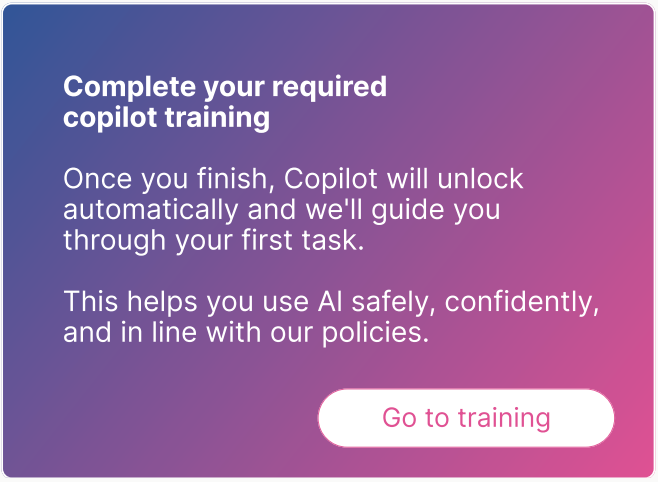

5. Context-aware misuse guidance

AI introduces powerful capabilities—but also meaningful risk. Nexthink protects employees in real time with contextual coaching that detects risky prompts as they are being written. When users inadvertently attempt to submit sensitive information or phrasing that violates policy, Nexthink intervenes instantly with safer, compliant alternatives. This prevents accidental data exposure, reinforces internal standards, and educates employees without slowing their workflow. By embedding compliance into the moment of action, organizations reduce risk while empowering employees to work confidently. Instead of relying on after-the-fact monitoring, Nexthink helps teams prevent issues before they occur, shaping safer, more responsible AI behavior.

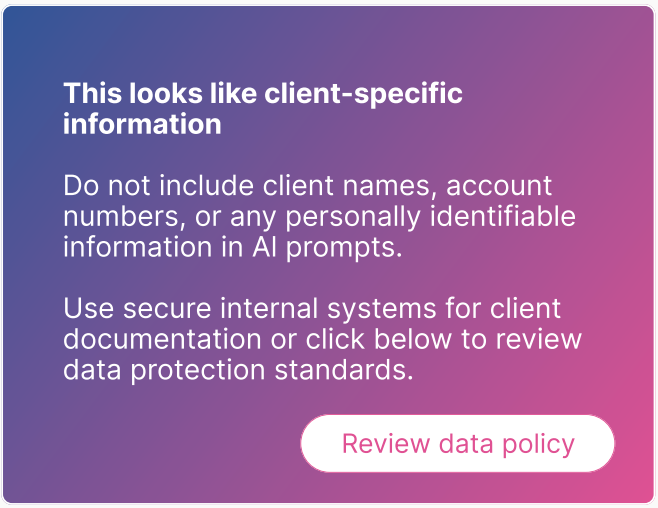

6. Staying compliant without slowing down

Employees often hesitate to upload data into AI tools because they fear unintentionally breaching policy. Nexthink solves this by delivering proactive guardrails just before a risky action occurs. Users see an immediate, clear warning outlining what cannot be shared—client details, proprietary data, legal restrictions—along with an acknowledgment step that reinforces safe behavior. This protects the organization while keeping work flowing smoothly, ensuring employees don't feel blocked or burdened. With compliance embedded directly into their workflow, staff can work faster and more securely. Nexthink gives teams confidence to use AI productively while eliminating the uncertainty that often prevents adoption.

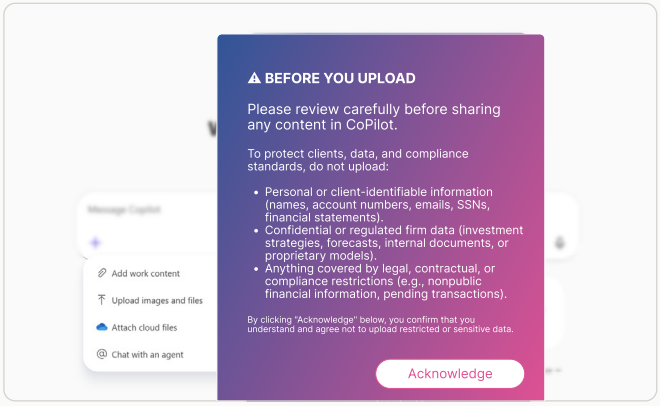

7. Provisioning access to advanced features

Access gaps are a major source of adoption friction: an employee attempts to use Copilot only to discover they lack the right license or permission. Nexthink removes this frustration through seamless automation. When the platform detects that a user needs an advanced Copilot capability, it can automatically trigger approvals, request necessary entitlements, or provision licenses immediately—no ticket, no delays, no manual intervention. This ensures employees always have the tools they need at the moment of intent, accelerating adoption and removing blockers that traditionally stall momentum. With Nexthink, access becomes effortless, and AI tools become instantly usable and scalable.

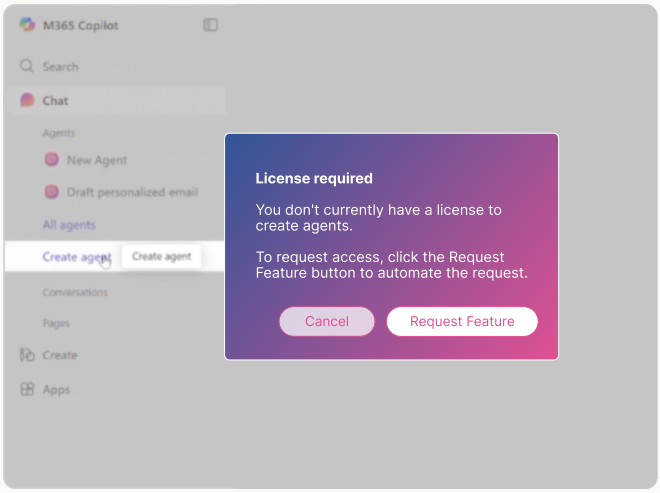

8. Unlocking better results with prompts

Employees often abandon AI tools because their outputs feel weak or unusable. Nexthink prevents this frustration by providing on-screen coaching that helps users improve their prompts in real time. Clear, contextual guidance shows employees how to refine phrasing, sharpen instructions, and adopt higher-quality prompting practices. Over time, this not only improves the usefulness of AI outputs but also builds employee confidence and skill. Rather than relying on trial and error, users get consistent, role-appropriate prompting support that makes AI more effective. Nexthink turns underwhelming results into a reliable, repeatable workflow that accelerates productivity across the enterprise.

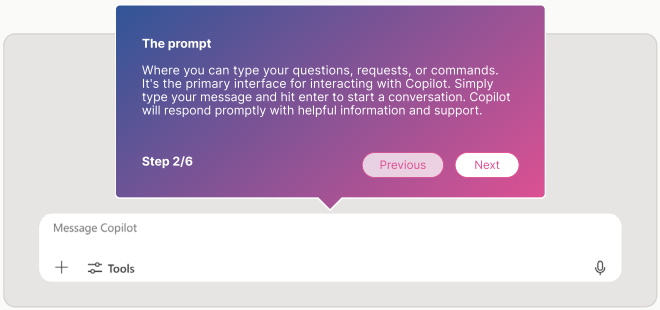

9. Instant policy updates for a rapidly evolving AI landscape

Policies around AI shift quickly as new risks, regulations, and best practices emerge. Nexthink ensures every employee stays informed by delivering instant updates directly inside the tools they use. When a compliance standard changes, all Copilot users receive targeted notifications, updated guidelines, and acknowledgment prompts—right within their workflow. This eliminates communication gaps and ensures policy adoption is not left to mass emails or intranet posts that go unnoticed. Leaders gain confidence that employees are aligned with the latest rules, and employees gain clarity without needing to seek information elsewhere. Nexthink keeps organizations compliant, consistent, and future-ready.

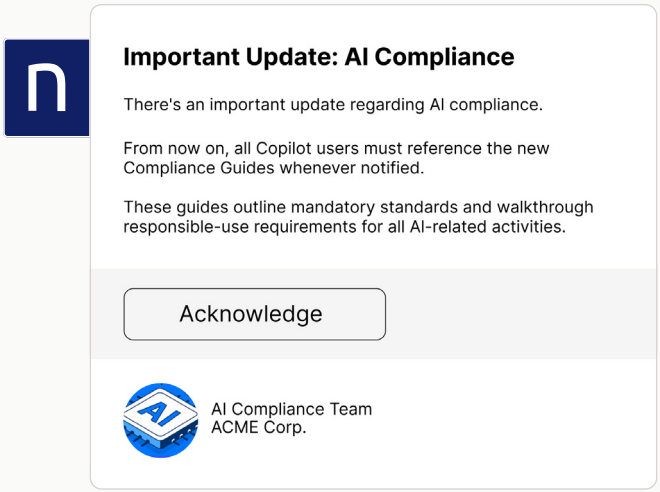

04.
Why only Nexthink is uniquely positioned to deliver AI Activation

Most organizations try to scale AI using one of two incomplete approaches: a Digital Adoption Platform (DAP) to guide users, or an analytics tool to measure usage. Yet neither one—and not even the combination of the two—can deliver true AI Activation. Guidance without visibility results in generic interactions that don't address the real barriers employees face. Visibility without enablement merely points out problems while leaving teams powerless to correct them. And neither approach can account for the mechanical, environmental, and workflow-level issues that quietly undermine AI programs long before people ever interact with the model.

Nexthink is the only platform that unifies visibility, guidance, governance, sentiment, and device-level intelligence into a single continuous loop—one that diagnoses, activates, and reinforces AI adoption end-to-end.

Why a traditional DAP is not enough

Traditional DAPs play an important role in showing users how to complete tasks, but they operate inside the boundaries of a single application. As a result, they cannot deliver a complete AI Activation strategy because they lack:

- Cross-application workflow orchestration (e.g., redirecting users from unapproved AI tools to sanctioned ones)
- Desktop-level targeting that goes beyond URL rules or page detection
- Policy-aware guardrails such as upload warnings, misuse coaching, or governance interstitials
- Automated provisioning or access workflows to give employees the right AI capabilities at the right moment
- Sentiment-driven interventions that adapt to user confidence, challenges, and feedback

A DAP can help someone once they are inside an AI tool, but it cannot shape the journey that happens before they enter it, nor can it influence behavior across the rest of their digital environment. AI adoption is an enterprise-scale challenge, and guiding a single UI is not enough to change culture, confidence, or compliance.

Why analytics platforms alone fall short

Analytics-only tools provide insights but cannot change outcomes. They can show low adoption or highlight unusual patterns, but they have no mechanism to intervene in the moment employees hesitate, make mistakes, or drift toward shadow AI. They cannot provide guidance, correct risky prompts, reinforce new behaviors, or connect usage patterns to root causes.

Critically, data-only platforms cannot distinguish between human reluctance and technical failure. Many adoption gaps stem not from lack of willingness but from poor device performance, unstable network connections, browser misconfigurations, or application friction. When a report shows low Copilot usage in a department, the cause may be slow page rendering or overloaded machines—not lack of training. Without visibility into the digital environment itself, analytics tools misdiagnose problems and misdirect effort.

Why Nexthink is singularly capable of bringing everything together

Nexthink solves all these gaps because it operates across the entire digital landscape—browser, desktop, network, workflow, and human experience—powered by a depth of telemetry no other platform can match. The Nexthink Collector, running on every device, gives organizations an unparalleled view of application performance, network health, system stability, configuration drift, software versions, and behavioral pathways across AI and non-AI tools.

This rich context allows Nexthink to distinguish when adoption stalls due to confusion, when it fails due to friction, and when it collapses due to environmental or technical obstacles. It also empowers Nexthink to intervene precisely: redirecting users from unapproved AI tools the moment they open them, launching walkthroughs as soon as someone shows signs of struggle, reinforcing policies when behavior drifts, and orchestrating access or permissions automatically when users need new capabilities.

Unlike standalone DAPs or analytics platforms, Nexthink links every component of AI adoption into one motion—usage, guidance, identity, device health, sentiment, governance, and outcomes—ensuring that insights immediately translate into action and that action is continuously measured for impact. It is the only ecosystem capable of diagnosing adoption barriers, activating behavior change, and proving value in a single, end-to-end loop.

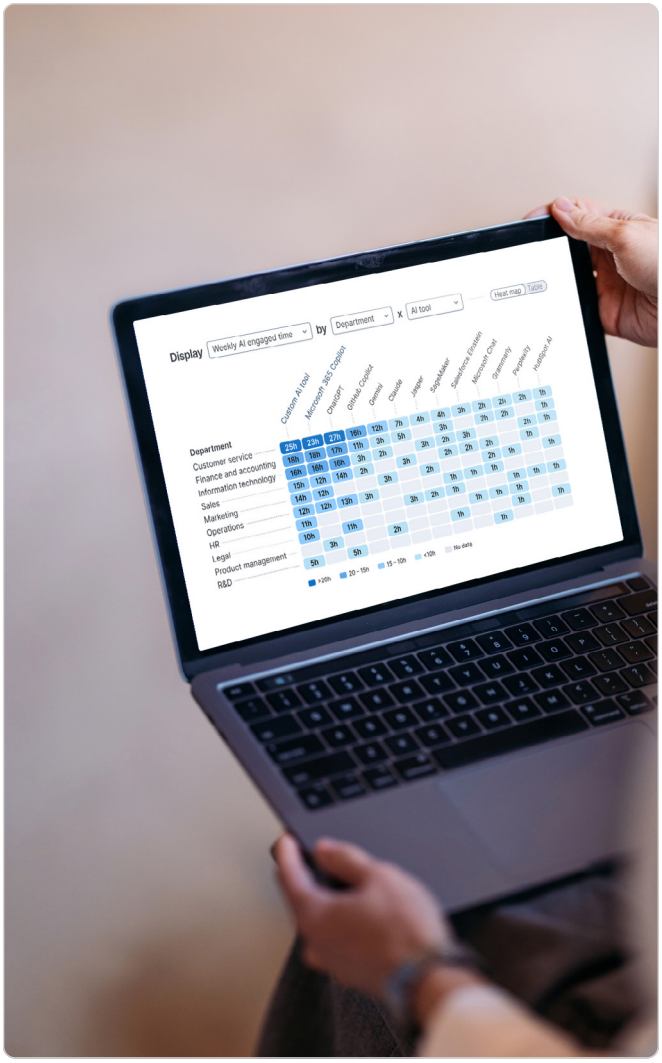

> One platform, one data model, one continuous loop, that's the power of Nexthink. With AI Drive and Adopt working together, everything connects: usage, guidance, identity, and outcomes. Organizations finally get a clear, unified view of their AI adoption and the precision to act on it confidently."

Guillaume Charles
Senior Director of Product Management, Nexthink

Conclusion

Enterprises don't struggle with AI because the models are lacking. They struggle because deploying AI without an adoption plan almost always leads to stalled pilots, shadow usage, compliance anxiety, and uneven capability across the workforce. The gap between "buying AI" and "embedding AI into everyday work" is where most initiatives fail—and where a clear activation strategy becomes essential.

Nexthink closes that gap by turning AI adoption into a true operational engine. AI Drive provides the visibility that organizations have historically lacked… revealing who is using AI, where value is emerging, where risk is accumulating, and where adoption is silently collapsing. Adopt meets those insights with in-flow guidance, steering users away from unapproved tools, coaching them through real tasks, reinforcing policy, and building confidence in the moment of need. Together, AI Drive and Adopt power a continuous loop that diagnoses friction, activates behavior change, and measures real outcomes across the enterprise.

With Nexthink, organizations no longer rely on manual campaigns or one-size-fits-all training. They can precisely target the right audiences—non-users, lagging departments, at-risk segments—and guide them with contextual experiences that raise confidence and close skill gaps. They can measure real lift and ROI, refine patterns based on live data, and build governance into daily workflows rather than bolting it on top.

For employees, this means AI becomes part of how work gets done, embedded in everyday processes and supporting smarter, faster decisions. For leaders, it means having the clarity and confidence to scale what works, course-correct what doesn't, and evolve their AI program continuously.

And because Nexthink runs on one platform, one data model, and one behavioral engine, this intelligence extends naturally beyond AI into any business application. Visibility becomes action… action becomes insight… and the organization moves forward in a continuous, self-improving loop.

Request a demo: nexthink.com/adopt

About Nexthink

About Nexthink is the global leader in Digital Employee Experience management. The company's products allow enterprises to create highly productive digital workplaces for their employees by delivering optimal end-user experiences. Through a unique combination of real-time analytics, automation and employee feedback across all endpoints, Nexthink helps IT teams meet the needs of the modern digital workplace.

nexthink

nexthink

Want to learn more about how Nexthink can help you save costs and improve employee experience?

Contact us