

The SME leader's playbook for practical AI applications

Reddal Insights — 7 January 2026

Kun Cao, Per Stenius

Artificial intelligence (AI) adoption for SMEs is moving beyond experimentation toward structured and scalable capabilities. The question is no longer whether to adopt AI, but how to make it deliver practical results. The article is available in English and Korean.

Artificial intelligence has shifted from novelty to real-time online infrastructure in the form of AI applications, such as ChatGPT, CoPilot and Gemini, which are built on top of large language models. The same assistants that once felt experimental are now embedded in email, documents, browsers, and line-of-business software. For small and midsize enterprises, the strategic question is no longer whether AI belongs to firm operations; it is how to adopt it in a way that compounds result instead of producing pilots that stall.

This article tries to answer those questions in a direct, operational way. Below are some points to guide you when reading this article:

- If the company is at the very beginning (level 0), start with personal experimentation and off-the-shelf tools. They will help identify low-risk, high-value use cases without waiting for major investments.
- If the team already uses AI applications informally, focus on the sections about upgrading to business-grade tools, defining simple policies, and running focused pilots (levels 1-2). These will help move from scattered experiments to managed impact.
- If the company is further along (levels 3-4), treating AI as a strategic capability, pay attention to the later chapters on governance, skills, and building an AI roadmap. They will help avoid the common pattern of “pilot fatigue” where enthusiasm fades after a few isolated successes.

Level 0: Personal experimentation to build intuition

Most organizations try to jump straight to platforms and vendors. That is premature. The first

job is to give leaders and managers a sense of what AI applications and tools can and cannot do. Treat the tools like you once treated spreadsheets: use them personally, not by proxy. Begin with low-stakes tasks such as practical research, rewriting emails, summarizing public articles, outlining a meeting agenda, and regularly practicing these habits that will carry through the journey. Ask clearly what you want, including the audience, tone, and constraints (in AI jargon, create a good 'prompt' that allows AI give you precise answers). Check with discipline, be mindful about privacy and confidentiality avoiding using sensitive information in open solutions, verifying facts, and reading for logic. It is better to iterate rather than accept the initial draft, because AI can build on previous prompts and improve through iterative requests. When a critical mass of managers has used AI for real work, can name where it helps and where it fails, and understands basic boundaries on confidential data, the organization is ready for structured use. However, we want to emphasize a "trust and verify" mindset – AI does make errors, and in this initial phase it is important to build understanding of its strengths and weaknesses by checking the answers and results it gives.

Level 1: Use off-the-shelf tools in daily work

The quickest gains come from applying general-purpose assistants and AI features already present in your software. This is not about custom systems. It is about making everyday work faster, clearer, and more consistent. Focus on recurring, text-heavy tasks in go-to-market, customer support, operations, finance, and HR. However, verify privacy and data security before uploading any sensitive materials or inputs to the tools.

- In sales and marketing, use AI to draft and localize copy, assemble proposals from a curated library of approved paragraphs, and prepare concise prospect briefs; measure the outcome in turnaround time and win rate.
- In support, let AI draft replies to common inquiries and summarize long threads while agents remain accountable for tone and truth; watch first-response time and resolution without escalation.
- In operations, convert workshop notes into standard operating procedures, extract checklists for handoffs, and publish meeting actions that name owners and dates; track time from decision to documented practice and rework incidents linked to unclear steps.
- In finance, turn numbers into narrative for management reviews and scenario memos, while keeping verification non-negotiable.
- In HR, accelerate job descriptions, interview kits, policy drafts, and staff FAQs, then judge success by time-to-post, early candidate quality, and fewer policy clarifications.

Adoption sticks when AI is used where work already happens. Keep the data boundary conservative, keep humans accountable for outputs, and measure impact with a few hard numbers rather than anecdotes. When those habits take root, AI shifts from curiosity to capability, and the firm is ready for the next levels. During this phase, turning recurring tasks into repeatable patterns creates most of the value in forms of prompts and templates. A good prompt does three things at once: it tells the tool who it should be, what outcome is required for which audience, and how to format the result so a human can verify and publish it quickly. A good template wraps that prompt with firm's voice, facts, and constraints so different people can produce work that looks like it came from a single individual^[1].

Level 2: Upgrade to secure, higher performance usage

Unmanaged enthusiasm creates hidden risks. Drafts with customer details slip into consumer tools, quality drifts as some teams accept outputs at face value, and tool sprawl overwhelms training and support. Ensuring larger scale AI adoption is done professionally starts by standardizing a compact, business-grade stack: choose what are the purposes where AI tools are needed and standardize them. The stack could contain one productivity suite with AI such as Microsoft365 with Co-Pilot or Google Workspace with Gemini, one general assistant for open-ended work such as ChatGPT or Perplexity, and a small number of specialized tools such as Claude for coding, where value is already proven^[2]. Insist on administrative control to ensure security and privacy, including single sign-on, multi-factor authentication, data-use controls that keep prompts and outputs out of model training, configurable retention, audit logs, and explicit data residency. Based on these approved, business grade AI tools, users can conduct safe experiments with controlled sandbox, so curiosity survives without compromising security.

Policy becomes practical when it is short, specific, and visible in the tools people use. Draw a clear line around prohibited data in general-purpose systems, state that the individual, not AI, remains responsible for accuracy and tone, set verification norms for claims and figures, define when and how to disclose AI assistance to clients, and restrict plugins and connectors to an allow-list. Technology and policy do not change behavior by themselves, so assign light-touch roles. An executive sponsor clears roadblocks, functional owners curate prompts and templates, and own outcomes; respected practitioners serve as champions who run clinics and surface edge cases, and a security lead approves tools and manages exceptions. Training shifts from generic introductions to hands-on sessions that use the firm's own documents and workflows, followed by quick clinics to fix what breaks. Avoid the predictable failure modes such as policy theater, which just causes over-restriction that fuels shadow IT, prompt sprawl, tool bloat, and automation without ownership. The trick is keeping the stack compact, embedding guidance where work happens, and assigning clear accountability.

You will know level 2 has been firmly reached when most target users complete named tasks with approved tools, guardrails are followed with only minor exceptions, a living library of prompts and templates is used, and leaders can point to several workflows with measurable performance gains. At that point, the constraint is a proxy to your own data and processes, not hygiene.

Level 3: Build internal capability

Level 3 implies moving on from “using AI” to “being good at AI.” Capability begins when AI gets closer to the information and workflows that define your business. The data job is not a grand transformation; it is making key facts about customers, products, orders, inventory, and projects clean enough and accessible enough to power a few priority use cases. Pick the systems of record that matter, agree on “good enough” definitions for critical fields, consolidate scattered spreadsheets into shared structures, and assign ownership for upkeep. If you cannot answer a basic question about a customer or a transaction without hunting through mailboxes and private folders, fix that first.

Process clarity is the multiplier. AI amplifies whatever it touches; if the workflow is ill-defined, you will scale confusion. Map the processes that matter such as lead to revenue, order to cash, ticket to resolution, and write down how they run. Then choose to use cases where text, judgment, and data meet in high-frequency work. A proposal assistant that assembles drafts from your own library, a project companion that turns notes and time entries into weekly reports, or a service helper that drafts responses based on your prior resolutions are all credible starting points. It's critical to tie each initiative with clear targets for impact, and stop or pivot if it does not create sufficient measurable values.

Create a small delivery team that combines a business owner empowered to make trade-offs, a technically fluent builder who can stitch tools and light automations together, and a respected frontline representative who keeps the design honest. Give them time, a visible backlog, and acceptance criteria. Run short pilots on narrow slices of the workflow, compare performance to the old way with numbers and qualitative feedback, and schedule iterations – the first version invariably will be clumsy. If a pilot clears the bar, deployment is change management: embed steps in the official process, update the procedures, train the next cohort on the use case rather than generic features, and maintain a single place where prompts, templates, and examples live. Resist the urge to scale everywhere at once. Let a handful of well-run deployments create confidence and reusable assets, such as a document library, a shared glossary, and a standard connector. As capability accumulates, the culture changes: teams experience AI as something that makes work clearer and results stronger, leaders ask tighter questions about decisions, errors, and customer simplicity, and learning compounds.

Level 3 ends when the organization can repeatedly identify a high-value task, connect it to the right data and process, build small tools that help the people working, and deploy it safely into daily operations. At that point, you have capability, not a project.

Level 4: Scale, govern and continuously improve

Reaching level 4 is less about adding another tool and more about changing how business learns. Treat successful use cases like small internal products with owners, version histories, and service levels. Build a simple roadmap that aligns a handful of initiatives to the outcomes that matter over the next twelve to twenty-four months, and fund fewer, better efforts with explicit metrics for graduation, iteration, pivot, or exit. Keep governance light but firm. Product owners can change prompts, templates, and workflows within boundaries, while anything that touches customer data, pricing, or regulated content gets a quick review from the data and security lead. When models update or outputs drift from the firm's tone, document the incident, fix it, and take the lesson back into training and policy. Over time, this creates an institutional memory that prevents repeated mistakes and sustains confidence.

Continuous improvement depends on feedback moving faster than habit. Give everyday users a low-friction way to flag what works and what breaks, and ship small fixes quickly, so teams see improvements are done promptly. As the portfolio grows, invest in basic data stewardship: stable interfaces, shared glossaries, consistent field names, validation at entry,

and sensible archival rules. The point is fewer surprises when workflows cross functions or when a new use case draws on the same sources. Advantages take root when AI fluency becomes part of how the firm hires, promotes, and develops, when communities of practice share patterns in language non-experts can use, and when incentives reward teams that consistently improve a workflow.

Risk management matures in parallel. Test model and plugin updates against a small suite of critical prompts before production, define graceful degradation paths for high-impact workflows, and take a consistent stance on disclosure in external communications so that client-facing teams can explain the role of AI without undermining confidence. None of this needs to be heavy; all of it should be intentional.

The payoff appears quietly before it becomes obvious. Customers notice that proposals arrive sooner and feel more tailored, that issues are resolved with fewer hops, and that your teams are consistent across languages and channels. Employees notice that documentation matches how the work is actually done, that suggestions are acted on, and that learning the next tool feels like an extension of what they already know. Owners notice that headline metrics move in the right direction and that gains persist even when a manager leaves or a vendor retires a feature. In that state, AI is no longer a program to manage. It is infrastructure for learning, a competitive stance rivals cannot copy quickly.

Conclusion - Start now, step-by-step gets you to scaled AI tool adoption

For small or midsize enterprises, the choice is not between adopting AI or ignoring it; it is between learning deliberately now or competing later against firms that did. The path is neither mysterious nor grandiose. Start by building intuition at the individual level, make use more professional with a compact and secure stack, connect the first few workflows to own data and processes, and then scale what works with light, visible governance. Do that with discipline and you will see that the economics of the business begins to change: cycle times shrink, error rates fall, and managers spend more time on judgment and less on drudgery. Most importantly, the company should develop a habit of improvement that rivals cannot copy quickly because it lives in how people think and work, not just in the tools they use.

Korean translation of the article

중소기업은 AI 도입을 고려할 때

중소기업(SME)은 인공지능(AI) 도입을 고려할 때 신중하게 접근해야 합니다. AI 도입은 “AI가 내 업무를 도와줄 수 있다”, “내 업무 효율을 높일 수 있다”라는 기대를 불러일으키지만, 실제 도입 시에는 예상치 못한 어려움이 발생할 수 있습니다.

중소기업은 AI 도입을 고려할 때, ChatGPT, Copilot, Gemini와 같은 LLM 기반 AI 도구를 도입할 때 신중하게 접근해야 합니다. AI 도입은 “AI가 내 업무를 도와줄 수 있다”, “내 업무 효율을 높일 수 있다”라는 기대를 불러일으키지만, 실제 도입 시에는 예상치 못한 어려움이 발생할 수 있습니다. AI 도입을 고려할 때, ChatGPT, Copilot, Gemini와 같은 LLM 기반 AI 도구를 도입할 때 신중하게 접근해야 합니다. AI 도입은 “AI가 내 업무를 도와줄 수 있다”, “내 업무 효율을 높일 수 있다”라는 기대를 불러일으키지만, 실제 도입 시에는 예상치 못한 어려움이 발생할 수 있습니다.

በ 4 ዓመታት ውስጥ የሚኖሩት የሥራ ስራዎች ለተጨማሪ የሥራ ስራዎች ሊሰጡ ይችላሉ። በተጨማሪም የሥራ ስራዎች ለተጨማሪ የሥራ ስራዎች ሊሰጡ ይችላሉ። በ 12~24 ዓመታት ውስጥ የሥራ ስራዎች ለተጨማሪ የሥራ ስራዎች ሊሰጡ ይችላሉ። በተጨማሪም የሥራ ስራዎች ለተጨማሪ የሥራ ስራዎች ሊሰጡ ይችላሉ።

በሌሎች አገሮች ውስጥ የሥራ ስራዎች ለተጨማሪ የሥራ ስራዎች ሊሰጡ ይችላሉ። በተጨማሪም የሥራ ስራዎች ለተጨማሪ የሥራ ስራዎች ሊሰጡ ይችላሉ። በተጨማሪም የሥራ ስራዎች ለተጨማሪ የሥራ ስራዎች ሊሰጡ ይችላሉ። AI ለሥራ ስራዎች ለተጨማሪ የሥራ ስራዎች ሊሰጡ ይችላሉ።

በሌሎች አገሮች ውስጥ የሥራ ስራዎች ለተጨማሪ የሥራ ስራዎች ሊሰጡ ይችላሉ። በተጨማሪም የሥራ ስራዎች ለተጨማሪ የሥራ ስራዎች ሊሰጡ ይችላሉ። AI ለሥራ ስራዎች ለተጨማሪ የሥራ ስራዎች ሊሰጡ ይችላሉ።

በሌሎች አገሮች ውስጥ የሥራ ስራዎች ለተጨማሪ የሥራ ስራዎች ሊሰጡ ይችላሉ። በተጨማሪም የሥራ ስራዎች ለተጨማሪ የሥራ ስራዎች ሊሰጡ ይችላሉ። AI ለሥራ ስራዎች ለተጨማሪ የሥራ ስራዎች ሊሰጡ ይችላሉ።

ቁጥር: የሥራ ስራዎች ለተጨማሪ የሥራ ስራዎች ሊሰጡ ይችላሉ **AI** ለሥራ ስራዎች ለተጨማሪ የሥራ ስራዎች ሊሰጡ ይችላሉ

በሌሎች አገሮች ውስጥ የሥራ ስራዎች ለተጨማሪ የሥራ ስራዎች ሊሰጡ ይችላሉ። በተጨማሪም የሥራ ስራዎች ለተጨማሪ የሥራ ስራዎች ሊሰጡ ይችላሉ። AI ለሥራ ስራዎች ለተጨማሪ የሥራ ስራዎች ሊሰጡ ይችላሉ።

References

[1] OpenAI, App use cases and prompts, retrieved from <https://help.openai.com/en/articles/12084614-app-use-cases-and-prompts> on January 7, 2026.

[2] Google Cloud, 1001 real-world gen AI use cases from the world's leading organizations, retrieved from <https://cloud.google.com/transform/101-real-world-generative-ai-use-cases-from-industry-leaders> on January 7, 2026.

In a follow-up interview, Global Leaders Insights spoke with Kun Cao, Client Director at Reddal, on how SME leaders can move AI from experimentation to practical impact. The full interview is available at [The SME Leader’s Playbook for Practical AI](#).

The topic was later covered by Electronic Times under the article '[ET Opinion] AI Is not just

for the big players'. The article is available in Korean via Electronic Times: [\[ET\]AI](#)
 - .