signal65

Measuring Modern AI Desktops: Real Productivity Gains for Office Workers

The AMD Ryzen™ AI 7 PRO 450G Brings AI to the Desktop

Ryan Shrout

In partnership with:

AMD
together we advance_

signal65.com

Contents

3 **How Office Workers Benefit from AI-Capable Desktops**

4 **The Office Worker**

6 **Testing Configuration**

7 **Compiling the Real-World Measurements**

9 **The Hardware Story**

11 **Create a Document**

13 **Create Meeting Notes**

15 **Click-to-Do Slide Update**

17 **Generate Tasks from Emails**

19 **Send Email with Update**

21 **Implications for Office Workers with AI Desktops**

23 **Appendix: Our Testing Theory**

25 **Important Information About this Report**

26 **System Configurations**

How Office Workers Benefit from AI-Capable Desktops

The desktop has long been the workhorse of corporate productivity, but the advancement of the AI PC has previously left it behind modern laptops in terms of capabilities. New AMD Ryzen AI desktops integrate dedicated AI acceleration onto the same silicon that has powered standard office productivity for decades, opening a new class of capability for the on-premise PC. For office workers who spend the bulk of their day across email, document creation, meetings, and content review, the question is no longer when AI tools will reach the desktop, but how much the right hardware unlocks.

Building on our previous analyses of office workers and road warriors, this study focuses on a question those earlier papers did not address directly. Does the silicon underneath the AI workload actually matter, and by how much? To answer it, we ran five common office workflows on two AMD desktop systems, a current-generation Ryzen AI 7 PRO 450G desktop and a prior-generation Ryzen 7 5700G desktop. Each workflow was tested with AI assistance on both systems and with a fully manual approach. The comparison isolates two effects: the gain from AI versus manual work, and the gain from new silicon versus old silicon on the same AI workload.

The results show that AI tools deliver clear and measurable productivity gains for office workers on capable AMD desktops. Across all five workflows, the new Ryzen AI 7 PRO 450G desktop completed in 19.1 minutes a set of tasks that took 56.7 minutes to do manually, a 66% reduction in time on a common subset of a typical workday. Equally important, the new silicon ran the same AI workloads 1.8x faster than the prior-generation desktop, and on two of the four workflows tested on both systems, the prior-generation desktop was actually slower with AI than just doing the work manually. Modern silicon is not a marginal upgrade for AI workflows, it is required to make those workflows worth running at all.

The case for current-generation AI desktops is therefore both immediate and broader than it appears. Immediate, because the productivity gains shown here are available now on shipping hardware and shipping software. Broader, because the same silicon trends that enable today's workflows are accelerating, and software increasingly assumes a capable NPU and modern memory subsystem to deliver on-device AI experiences. Office workers running prior-generation desktops will increasingly find themselves limited not by their willingness to adopt AI tools but by the hardware required.

Key Highlights:

- The new AMD Ryzen AI 7 PRO 450G desktop **reduces office work time by 66%** in a typical AI-enabled workflow, completing in 19.1 minutes what manual workflows take 56.7 minutes to accomplish.
- Office workers using AI tools on a modern AMD desktop **save approximately 1.1 workdays per week** at typical task frequencies, the equivalent of roughly 11 work weeks per year.
- Individual workflow steps run as much as **17x faster with AI** than manual work, with meeting summarization tasks showing the largest gain.
- New silicon runs the same AI workload **1.8x faster than the prior-generation** AMD desktop, and on certain workflows old silicon makes AI assistance slower than the manual approach.

The Office Worker

The Office Worker persona covers professionals who spend the majority of their workday at a desktop computer, working across an integrated suite of business applications. This persona is the largest segment of the corporate workforce by population, and its productivity ceiling is set primarily by the desktop hardware, the speed of the AI tools that desktop can run, and the cognitive overhead of switching between applications. The persona includes:

- Knowledge workers in marketing, communications, sales operations, finance, project management, human resources, and other corporate functions
- Heavy users of document creation, email, presentations, and meeting collaboration tools across an integrated productivity suite
- Workers producing client-facing or internal content (briefs, presentations, status updates, meeting summaries) on a regular cadence
- Workers operating from a fixed office or home-office environment where the desktop is the primary work surface and its performance is the main constraint on responsiveness
- Workers whose individual productivity gains compound directly into team and organizational throughput

To measure the impact of AI assistance on this persona, we identified five workflows that together represent a meaningful slice of a typical office worker's day. The workflows span content drafting (Create a Document), summarization (Create Meeting Notes), multi-application UI interactions (Click-to-Do Slide Update), structured extraction (Generate Tasks from Emails), and tone-conscious email composition (Send Email with Update). Four of the five workflows use local large language models running on the desktop. Running these models locally on the desktop keeps the underlying data on the device, avoiding the cloud egress and licensing costs that come with cloud-only AI workflows. That security and cost advantage depends on hardware capable enough to deliver acceptable model latency. The right desktop is the first decision in a combined hardware-software-policy rollout, not just a hardware refresh.

The map of workflows to AI tools is shown on the following page.

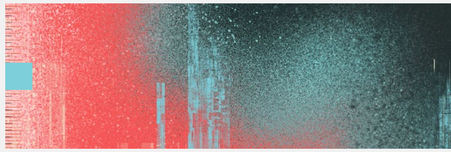

REAL PRODUCTIVITY GAINS FOR OFFICE WORKERS

The Office Worker

Create a Document	LM Studio	Quickly create a first draft to meet a deadline.
Create Meeting Notes	Generate AI PC Liquid	Create formatted AI notes from meeting transcription.
Click-to-Do Slide Update	Click-to-Do	Use Click to Do to quickly update a presentation.
Generate Tasks from Emails	AnythingLLM	Save time reading emails and planning daily schedule.
Send Email with Update	AnythingLLM	Change email tone from casual to formal.

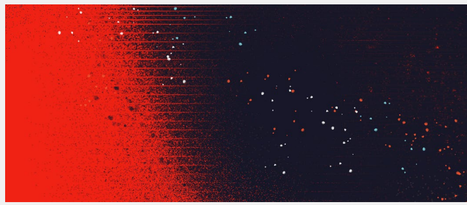

Testing Configuration

Our analysis uses two different AMD desktop systems to isolate two effects simultaneously. The first is the productivity gain from AI assistance over manual workflows. The second is the gain from current-generation silicon over prior-generation silicon on the same (or as close as we can get it) AI workload. The new system is an HP EliteDesk 8 SFF built on the AMD Ryzen AI 7 PRO 450G, which integrates a dedicated NPU alongside the CPU and GPU. The reference prior-generation system is an HP Pavilion built on the AMD Ryzen 7 5700G, a capable desktop CPU from a roughly 5-year-old platform generation that predates the integration of AI-specific silicon. Both systems were configured with 32GB of RAM, with the new system using faster DDR5 memory and the older system using DDR4. The manual workflow is hardware-independent and applies equally to either system in our experience, providing a stable baseline against which both AI configurations can be compared.

Each workflow was run three times on each applicable configuration and the results averaged. For local-LLM workflows, the same models, prompts, and inputs were used on both systems. The Click-to-Do workflow was tested only on the new system because the prior-generation Ryzen 7 5700G does not meet the Copilot+ PC requirement for an integrated NPU at 40 TOPS or above. This exclusion is itself a finding. Some AI experiences are gated by silicon class and are not available on older hardware at any speed.

	MODERN SYSTEM - HP ELITEDESK 8 SFF	PRIOR-GEN REFERENCE - HP PAVILION
CPU	AMD Ryzen AI 7 PRO 450G	AMD Ryzen 7 5700G
Integrated NPU	AMD XDNA 2	Not present
Graphics	AMD Radeon (integrated)	AMD Radeon (integrated)
RAM	32GB DDR5	32GB DDR4
Storage	1TB NVMe SSD	1TB NVMe SSD
Operating System	Windows 11 Pro 26200	Windows 11 Pro 26200
Copilot+ PC Eligible	Yes	No

Compiling the Real-World Measurements

Our testing shows that the new AMD desktop, paired with current AI tools, delivers measurable productivity gains across every workflow we measured. Across the complete Office Worker test suite, AI-enabled workflows on the Ryzen AI 7 PRO 450G system completed in approximately 19 minutes, compared to nearly 57 minutes for the same outcomes produced manually. That is a 66% reduction in total task time on a common subset of a typical workday, and a 3.0x speed-up across the five workflows. These gains are in no way marginal, and they are available today on shipping hardware and shipping software.

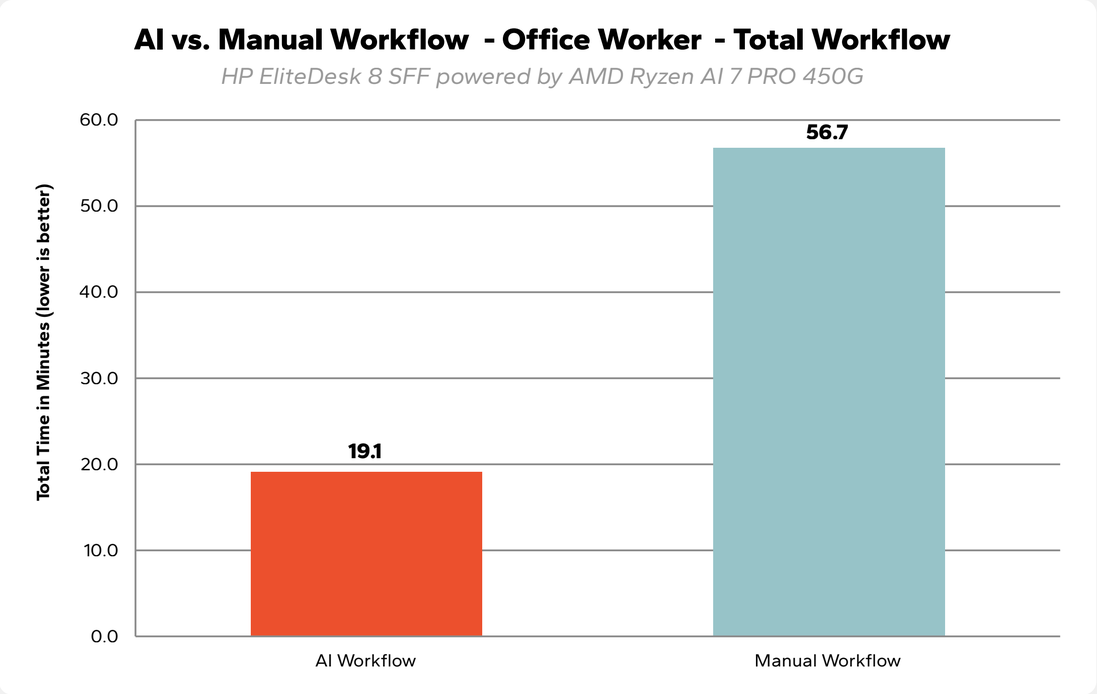

Projecting these gains to a typical workweek illustrates the impact at scale. At realistic task frequencies (one document creation, three meetings, two slides, three task batches, and ten status emails per day) the AI-enabled workflows save approximately 1.1 workdays per week, or roughly 11 work weeks per year. The savings are not evenly distributed across workflows. Meeting summarization alone accounts for more than two-thirds of the total weekly recovery because it combines a high task frequency with a workload that AI models are incredibly well-suited to handle. The other workflows show more modest but consistent gains in the 1.2x to 3.0x range, and those smaller gains still compound through daily repetition.

Compiling the Real-World Measurements

Workflow	Frequency	Per-Task Savings	Weekly Savings	Share of Total
Create a Document	1 / day	6.7 min	**33.6 min**	6%
Create Meeting Notes	3 / day	26.8 min	**401.9 min**	74%
Click-to-Do Slide Update	2 / day	1.6 min	**15.8 min**	3%
Generate Tasks from Emails	3 / day	0.9 min	**13.6 min**	2%
Send Email with Update	10 / day	1.6 min	**81.4 min**	15%
TOTAL			**546.3 min / week**	100%

Equivalent to approximately 9.1 hours per week, or 1.1 workdays per week recovered.

Beyond raw time savings, AI assistance changes the texture of office work in ways that compound over time. Summarization reduces cognitive load when reviewing meeting transcripts and email threads. Tone refinement reduces the polish-and-rewrite cycle on outbound communications. Structured extraction surfaces action items that would otherwise be missed in long correspondence. Multi-application AI features such as Click-to-Do reduce context switching, which has its own well-documented cost on knowledge worker output. The complete productivity story for AI-capable desktops is the sum of these effects, not just the raw seconds saved on each workflow.

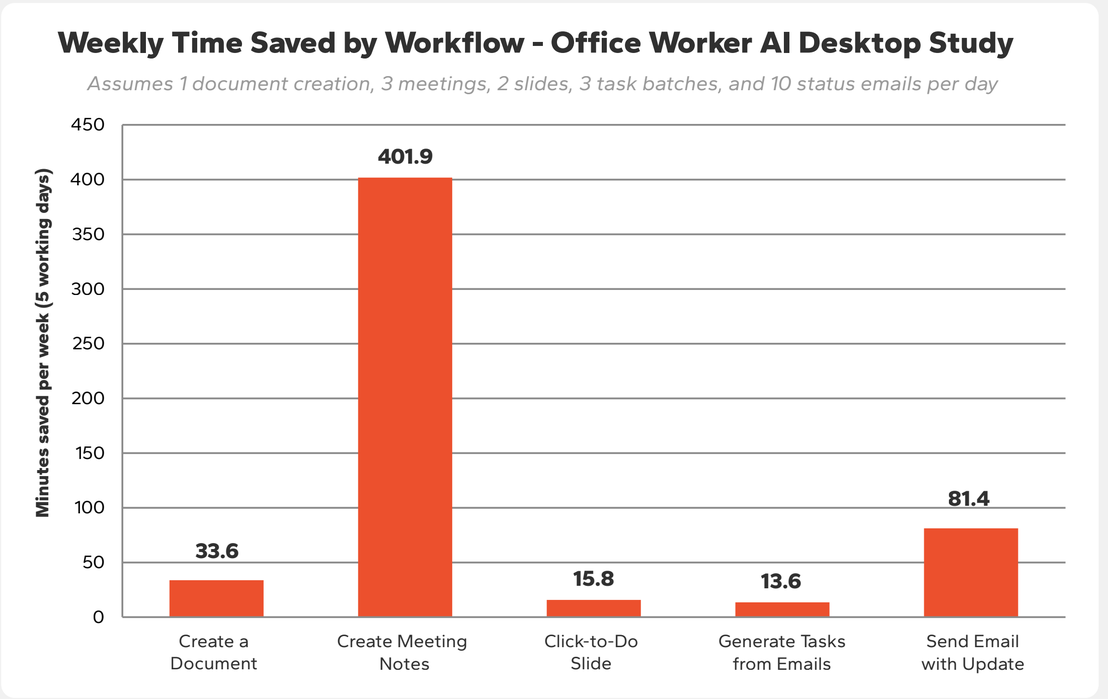

Office workers using AI tools on a modern AMD desktop **recover the equivalent of 1.1 workdays each week**, totaling roughly 11 work weeks of productive capacity per year without extending work hours or adding headcount.

The Hardware Story

On the prior-generation Ryzen 7 5700G desktop, two of four AI workflows take longer than the same work done by hand. The new Ryzen AI 7 PRO 450G turns every workflow into an AI win.

New Silicon is What Makes AI Workflows Worth Running on the Desktop

The single-system comparison above answers one important question. Does AI assistance save time on this hardware? Another important question for IT decision makers, however, is whether the silicon in the system actually matters for AI workloads, and if so by how much. To answer that, we ran the same AI workflows on a prior-generation AMD Ryzen 7 5700G desktop and compared the results to the new Ryzen AI 7 PRO 450G desktop. The manual workflow stays identical on both systems and serves as a stable reference baseline.

The results show that modern silicon does not just speed up AI work, it determines whether AI work is worth doing at all. Across the four workflows tested on both systems, the new Ryzen AI 7 PRO 450G desktop ran the same AI workloads on average 1.8x faster than the prior-generation system, recovering 14.3 minutes per workflow set from hardware alone. Two of those four workflows tell an even stronger story. On the prior-generation Ryzen 7 5700G, the AI workflow took longer than doing the same work by hand. Generate Tasks from Emails took 10.3 minutes on the older system with AI assistance, compared to 4.9 minutes manually. Send Email with Update took 4.3 minutes with AI on the old system, essentially identical to the 4.2 minutes required manually. In both cases, current silicon makes the difference between AI assistance being a productivity win and being a net loss.

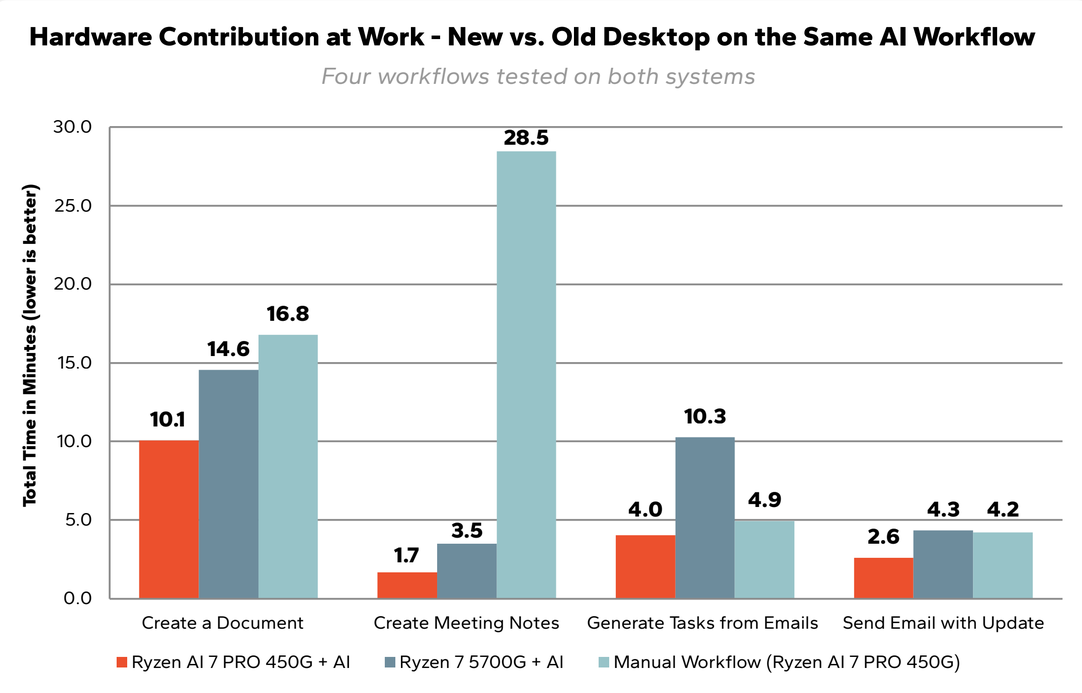

The Hardware Story

The underlying performance metrics make the source of the speed-up clear. On the local large language models used in this study, the new desktop delivers Time-to-First-Token between 3.3x and 5.3x faster than the older desktop, and generation throughput between 2.3x and 2.6x higher. Time-to-First-Token determines how long the user waits before the model begins producing output, which is the most psychologically important latency for interactive AI work. Generation throughput determines how quickly the rest of the response arrives. Together they explain why the new desktop closes the gap between AI and manual on the smaller workflows and dramatically widens it on the larger ones.

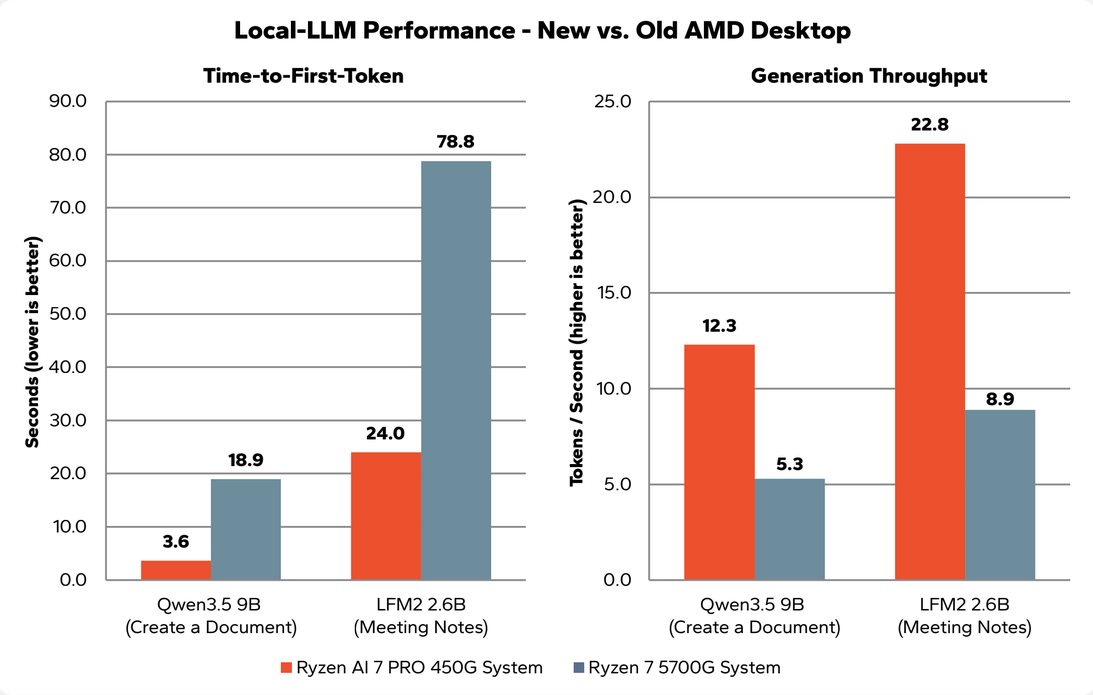

Hardware refresh planning for office workers in 2026 and beyond should treat AI capability as a primary specification, alongside memory, storage, and security baselines. Specifying desktop systems without integrated NPUs and modern memory subsystems will not just leave AI gains on the table. It will actively impair productivity for any worker whose workflows include AI-assisted steps.

LM Studio

Create a Document

Creating a new document is a common but demanding task for marketing, communications, and product teams. The work involves taking source material (a product brief, an internal announcement, or a set of structured bullet points) and producing a polished document. The Create a Document scenario tests how much of that initial drafting can be productively offloaded to a local AI model, while keeping the manual polish step where human judgment is most valuable.

The AI-enabled workflow uses a 9B-parameter Qwen3.5 local model running on the AMD desktop through LM Studio. The user supplies a structured prompt with the source material and the model produces a complete draft. The human then performs an identical polish step (adjusting brand voice, adding specific information and data, and final proofreading) on both the AI workflow and the manual workflow. The manual workflow simply skips the AI generation step and writes the document from scratch before the same polish.

AI-Enabled Steps	Manual Steps
The AI version uses a local large language model to generate a complete first draft from a structured prompt, which the writer then polishes.	The manual version writes the entire draft from scratch and applies the same polish step at the end.
Open source material and target document	Open source material and target document
Launch local LLM (Qwen3.5 9B) via LM Studio	Draft the document from scratch in Word
Submit structured prompt with source material	
Receive and review AI-generated draft	
Apply manual polish (brand voice, specifics, proofreading)	Apply manual polish (brand voice, specifics, proofreading)

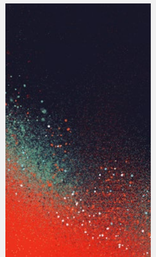

REAL PRODUCTIVITY GAINS
FOR OFFICE WORKERS

LM Studio — Create a Document

AI-assisted drafting **cuts document creation time by 40%** on a modern AMD desktop, freeing writers from the cognitive load of the blank page while preserving the human polish step that matters most.

Performance results show that the AI-enabled workflow completes the entire document creation process in approximately 10 minutes, compared to nearly 17 minutes for the manual approach, a 40% reduction in time and a 1.7x speed-up. The AI portion takes about 5 minutes of model inference and review, replacing roughly 12 minutes of manual drafting. The hardware-independent polish step (313 seconds in our testing) is identical on both sides, so the net AI savings come entirely from the drafting portion.

The quality of AI-drafted documents in our testing is suitable for downstream polish, capturing the core structure, key facts, and a serviceable first-pass voice. The result is not publishable without human refinement, but it consistently saves the writer the time and cognitive energy required to produce a clean first draft. For teams producing multiple documents per week, that drafting time compounds quickly, and the consistency of structure across drafts also reduces the editorial overhead on downstream reviews.

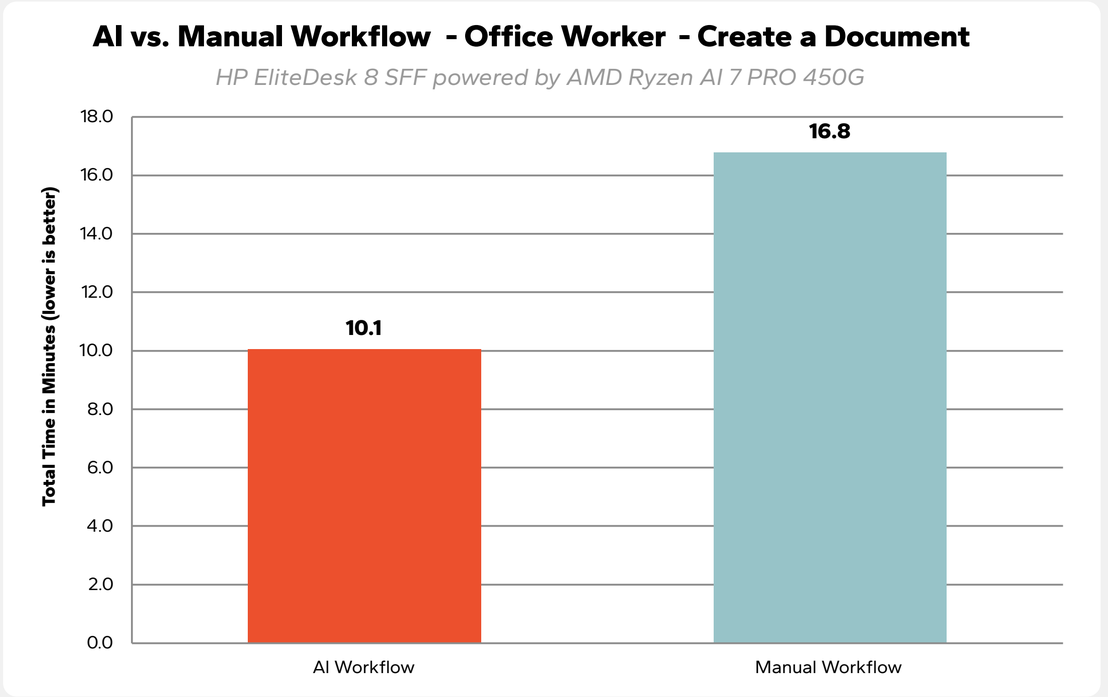

Workflow Total Time Savings (minutes)

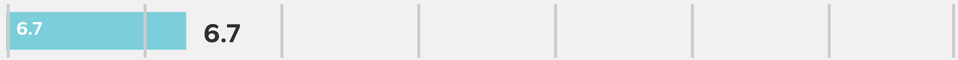

Create Meeting Notes

Meeting summarization is one of the most time-consuming recurring tasks in commercial knowledge work, performed multiple times per week by most office employees. It is also a task category that current AI models are particularly well-tuned for. The work involves taking a meeting transcript and producing a structured, readable summary with decisions, action items, and follow-ups. Manual production typically requires reviewing the transcript or recording, taking notes, and reformatting them into a shareable document.

Our test uses a recorded meeting transcript and asks an LFM2 2.6B local model to produce a structured summary with decisions, action items, and key discussion points. The manual workflow simulates a knowledge worker reviewing the same transcript and producing the same summary by hand. Both approaches target an identical deliverable: a clean meeting notes document suitable for distribution to stakeholders who did not attend.

AI-Enabled Steps	Manual Steps
The AI version uses a small local model specifically tuned for transcript summarization to produce structured meeting notes.	The manual version reads the same transcript and produces the same notes by hand, including formatting and cleanup.
Load meeting transcript	Load meeting transcript
Launch local LLM (LFM2 2.6B) via LM Studio	Read transcript carefully
Submit transcript with structured summary prompt	Take notes by hand
Receive AI-generated notes	Reformat notes into structured document
Light review and formatting	Final cleanup and formatting

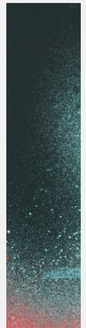

REAL PRODUCTIVITY GAINS FOR OFFICE WORKERS

Generate AI PC Liquid

Create Meeting Notes

Meeting summarization runs **17x faster with a local AI model** on a modern AMD desktop, recovering more than six hours per week for the typical office worker.

The results are the most eye-opening in this study by a wide margin. The AI-enabled workflow completes in approximately 1.7 minutes, compared to 28.5 minutes manually, a 94% reduction in time and a 17x speed-up. This gain sits well outside the 1.2x to 3.0x range seen on the other four workflows. The reason is that summarization plays directly to two strengths of modern small models, structure extraction and rewording from long input passages. The same characteristics that limit small models on open-ended reasoning tasks (limited reasoning depth, occasional factual drift) matter much less for summarizing structured conversational text where the source material is already present.

For an office worker attending three meetings per day, the cumulative time saved on note creation alone is approximately 80 minutes per day, or over 6 hours per week. Meeting summarization is by itself a credible justification for the AI-capable desktop refresh, even before the other workflows are considered. Beyond raw time savings, AI-generated notes are more consistent in structure and less prone to gaps caused by participant attention during the meeting itself, since the summarization step happens against a complete transcript rather than partial real-time notetaking.

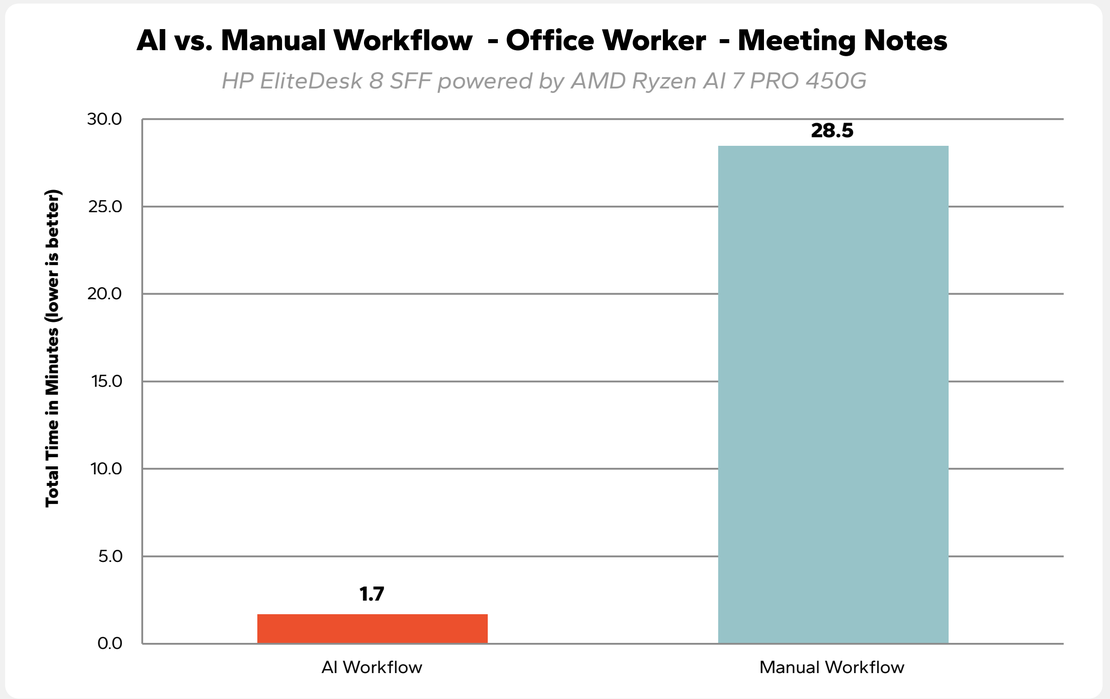

Workflow Total Time Savings (minutes)

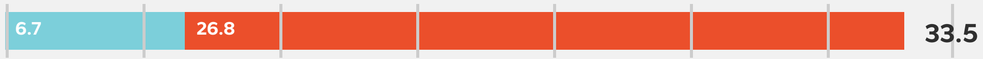

Click-to-Do

Click-to-Do Slide Update

Updating a slide with refreshed content represents a common multi-application workflow that requires coordinating information across PowerPoint, a source document, and image editing. Our Click-to-Do scenario tests Windows Click-to-Do, a Copilot+ PC feature that allows the user to point at on-screen content and invoke AI actions directly without switching applications.

The AI workflow uses Click-to-Do to select an image in PowerPoint and apply background blur, then selects source text in a Word document and instructs Click-to-Do to summarize it into bullet points, which are then pasted into the slide. The manual workflow performs the same operations the traditional way, switching between PowerPoint, Paint, and Word, and copying and editing each element by hand.

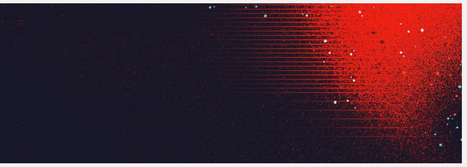

AI-Enabled Steps	Manual Steps
The AI version uses Click-to-Do, a Windows Copilot+ feature, to act on on-screen content across applications without manual context switching.	The manual version performs the same operations by switching applications and editing each element by hand.
Open PowerPoint deck and source Word document	Open PowerPoint deck and source Word document
Use Click-to-Do to select image in slide, apply background blur	Copy image from PowerPoint, switch to Paint to edit, paste back
Use Click-to-Do to select text in Word doc, summarize into bullets	Switch to Word, read and mentally summarize content
Paste AI-generated bullets into PowerPoint	Switch back to PowerPoint, type summary into text box
Light formatting cleanup	Adjust formatting manually

REAL PRODUCTIVITY GAINS FOR OFFICE WORKERS

Click-to-Do

Click-to-Do Slide Update

Click-to-Do **completes multi-application slide updates 3.0x faster** on a Copilot+ AMD desktop while cutting context switches from seven to two.

The AI workflow completes in 47 seconds versus 141 seconds for the manual approach, a 67% reduction in time and a 3.0x speed-up. In addition to the raw time savings, the AI workflow reduces context switches from seven application transitions to two, which reduces cognitive disruption. Multi-application workflows benefit twice from AI assistance, once from automation of the work itself, and again from staying in a single visual context.

It is worth noting that Click-to-Do is gated by the Copilot+ PC requirements, which include integrated NPU performance of at least 40 TOPS. The prior-generation Ryzen 7 5700G system tested in this study does not meet those requirements, so this workflow was tested only on the new Ryzen AI 7 PRO 450G desktop. This is a useful data point in itself; some AI experiences are not available at all without modern silicon, and are not addressable through any amount of cloud-side processing because the experience itself depends on the on-device feature being present.

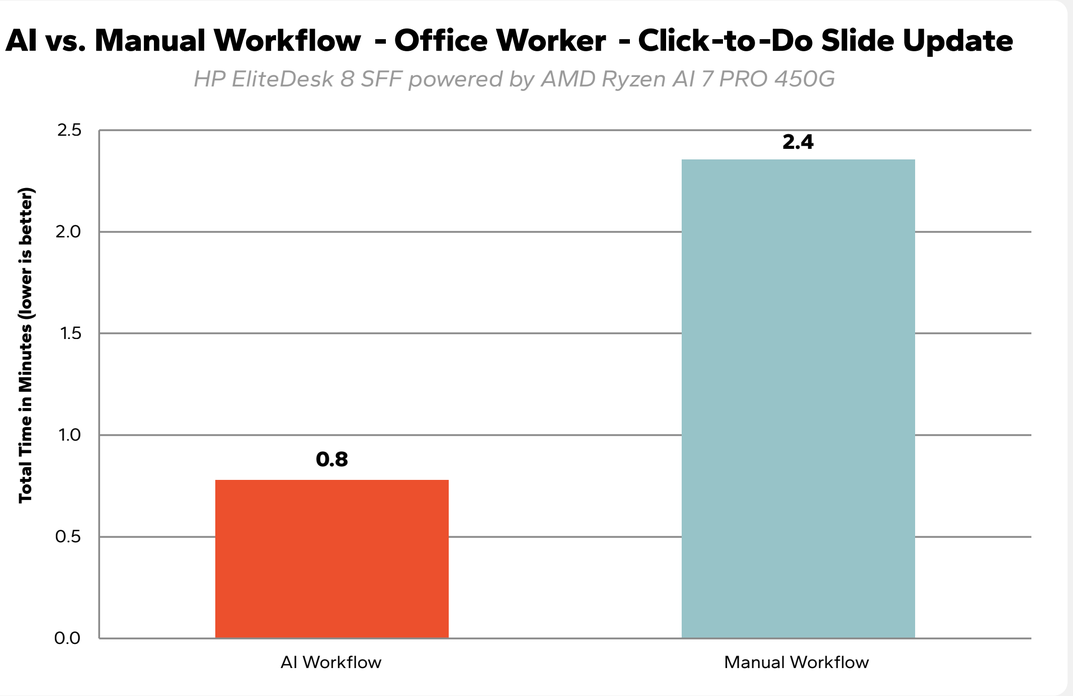

Generate Tasks from Emails

Knowledge workers regularly need to extract actionable items from a thread or batch of emails. The work involves reading through messages, identifying what requires action, and producing a structured list of tasks with owners and deadlines where possible. Our test asks a Qwen3.5 4B local model to perform that extraction from a sample email batch, compared to a knowledge worker doing the same work by hand.

This is also the workflow where the AI versus manual gap is smallest, and that is itself an instructive result. The source material is short enough that an experienced reader can extract the same task list relatively quickly, and the model inference adds latency that partially offsets the reading savings. On capable silicon, the AI workflow still wins, but the margin is modest. On prior-generation silicon, the margin collapses entirely.

AI-Enabled Steps	Manual Steps
The AI version uses a small local model to extract structured task items from a batch of emails.	The manual version reads each email and extracts task items by hand.
Load email batch into the workflow	Load email batch into the workflow
Launch local LLM (Qwen3.5 4B) via AnythingLLM with Lemonade Server	Read each message carefully
Submit emails with task-extraction prompt	Identify task items as they appear
Receive structured task list	Compile structured task list
Light verification of extracted tasks	Review the list for completeness

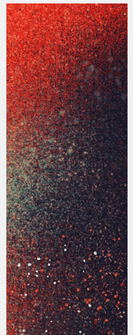

AnythingLLM

Generate Tasks from Emails

AI task extraction is modestly faster on a new AMD desktop and slower than manual on a prior-generation desktop. The same workload, the same model, the same input, only the silicon differs.

On the new Ryzen AI 7 PRO 450G desktop, the AI workflow completes in approximately 4 minutes versus 4.9 minutes manually, a modest 18% improvement and a 1.2x speed-up. The savings are real but smaller than the most dramatic workflows in this study. The same workflow on the prior-generation Ryzen 7 5700G, however, takes 10.3 minutes (see chart on page 9), which is 2.1x slower than just doing the work by hand. This is the clearest example in our testing of why hardware matters for AI workflows. On capable silicon the workflow is a modest win, and on prior-generation silicon it is an active loss.

Office workers do not benefit from running AI on hardware that cannot keep up with the latency requirements of the task. Where AI inference takes longer than the manual work it replaces, the AI option is simply worse, no matter how capable the underlying model. This is the reason the new silicon matters even on workflows where the absolute time savings appear small. Without it, the workflow does not save time at all.

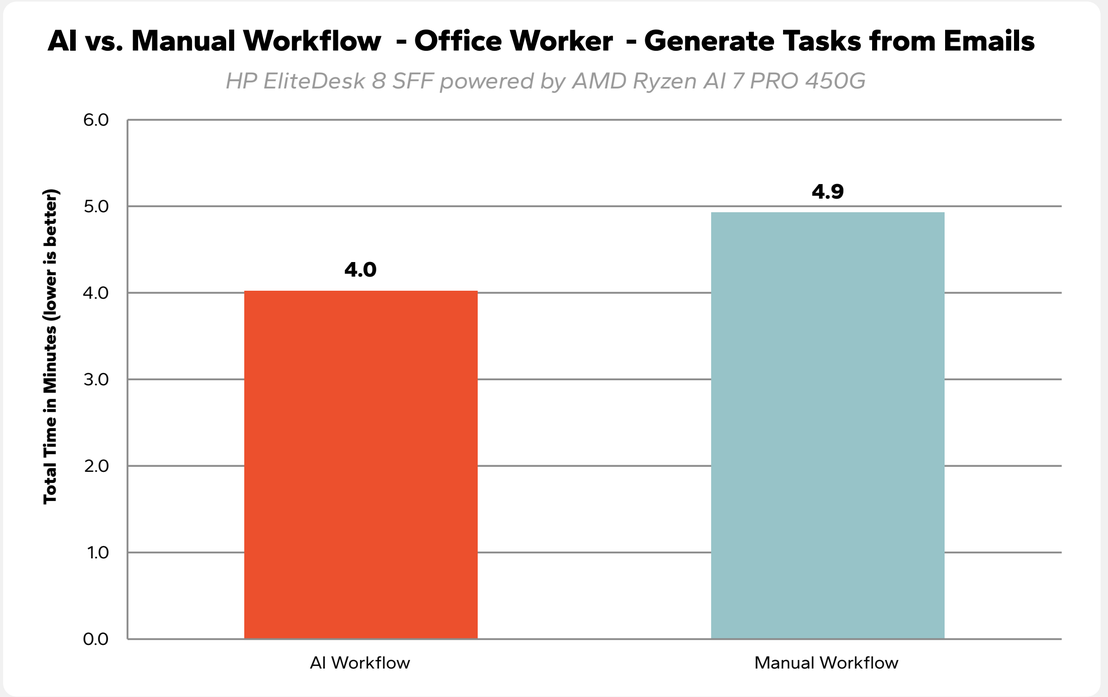

Workflow Total Time Savings (minutes)

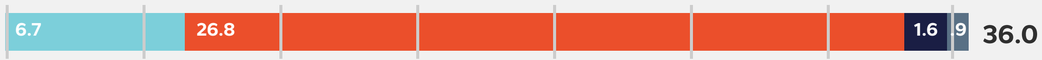

AnythingLLM

Send Email with Update

Composing professional status emails is a high-frequency task across roles. Workers commonly write a casual draft and then refine it for tone, structure, and audience appropriateness, or they assemble a status update from a set of bullet points. Our Send Email with Update scenario tests how much of that refinement step can be offloaded to a local AI model.

The AI workflow uses a Qwen3.5 4B local model to draft and tone-adjust the message based on bullet inputs from the writer. The manual workflow has the writer compose the same message from scratch and refine it for tone and formatting by hand.

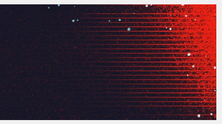

AI-Enabled Steps	Manual Steps
The AI version drafts and tone-adjusts a status email from bullet inputs using a small local model.	The manual version writes the email from scratch and refines tone by hand.
Compose informal bullet-point notes for the update	Compose informal bullet-point notes for the update
Launch local LLM (Qwen3.5 4B) in AnythingLLM	Write the email from scratch
Submit notes with tone and format prompt	Refine tone and formatting manually
Receive AI-drafted formal email	
Review and send	Review and send

REAL PRODUCTIVITY GAINS
FOR OFFICE WORKERS

AnythingLLM

Send Email with Update

AI-assisted email composition **saves 38% per message** on a modern AMD desktop. The same workflow on a prior-generation desktop provides no meaningful advantage over writing the email manually.

The results show a 38% reduction in time on the new desktop, with the AI workflow completing in 2.6 minutes versus 4.2 minutes manually, a 1.6x speed-up. The per-incident savings are smaller than the most dramatic workflows in this study, but they compound substantially at typical task frequencies. At ten status emails per day, the AI workflow recovers more than 16 minutes daily, or over an hour per week, on this task alone.

The contrast with the prior-generation system is again instructive. On the AMD Ryzen 7 5700G, the AI workflow takes 4.3 minutes, essentially identical to the 4.2-minute manual time. There is no meaningful productivity advantage to running AI for this task on the older system. The new silicon recovers the AI advantage entirely, and turns a wash into a clear win.

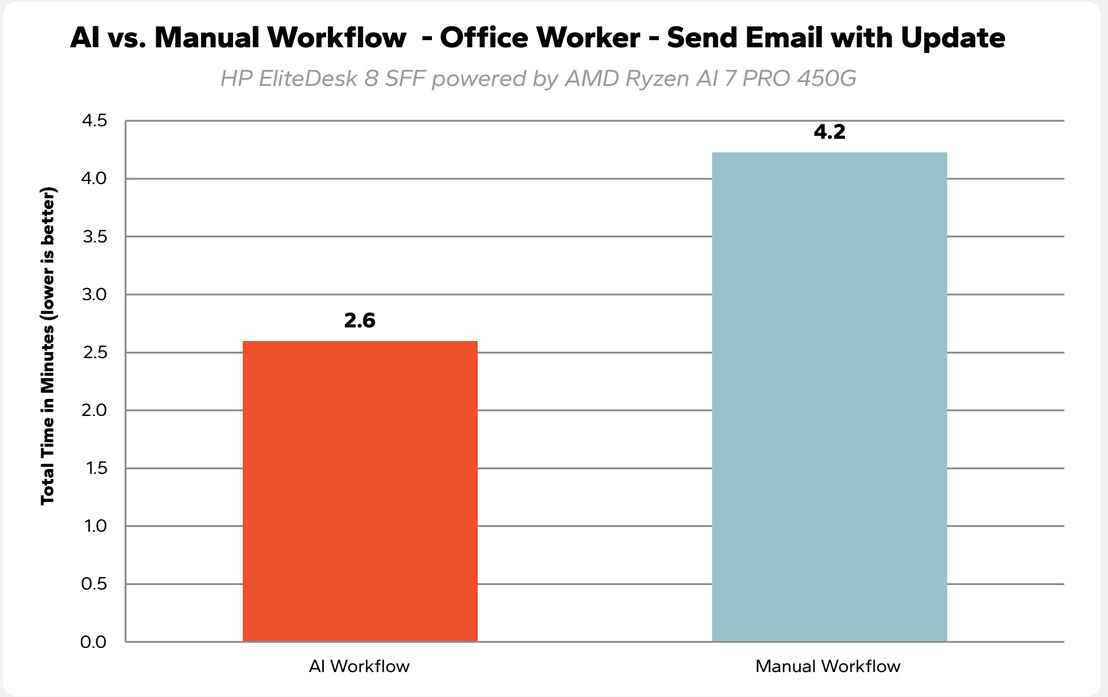

Workflow Total Time Savings (minutes)

Implications for Office Workers with AI Desktops

Our testing shows that AI-capable AMD desktops deliver clear and measurable productivity advantages for office workers. Across real-world workflows spanning content drafting, meeting summarization, multi-application UI interactions, structured extraction, and tone-conscious composition, AI-enabled workflows on the new AMD Ryzen AI 7 PRO 450G desktop completed in 19.1 minutes a set of tasks that took 56.7 minutes manually, a 66% reduction. For office workers whose daily output is constrained by the speed and reliability of their primary work surface, that difference represents a meaningful expansion of weekly productive capacity. The headline figure, approximately 1.1 workdays recovered per week, is the largest single productivity gain we have measured in any of our AI desktop studies to date.

The unique contribution of this study, however, is the direct measurement of silicon contribution on the same AI workload. By running identical AI workflows on a current-generation Ryzen AI 7 PRO 450G desktop and a prior-generation Ryzen 7 5700G desktop, we isolated the hardware effect at 80% on average across the four workflows where both systems were tested. More striking, in two of those four workflows, the prior-generation desktop was actively slower with AI than doing the work by hand. The lesson is clear. AI is not a feature that benefits from being added to any desktop. It is a workload class with its own hardware requirements, and prior-generation systems will increasingly fail to meet them as software assumes capable NPUs and modern memory subsystems.

For IT and business decision makers, the implication is a refresh planning question rather than an adoption question. The productivity gains shown here are available now, but only on hardware specified for them. Refresh cycles that do not prioritize AI silicon will leave the largest single productivity opportunity of the current desktop generation on the table. Refresh cycles that do prioritize AI silicon, on the other hand, can justify the upgrade on AI workflow gains alone, before factoring in the broader benefits of newer platforms (improved security baselines, faster memory and storage subsystems, better thermal envelopes, longer support windows).

From an IT decision-maker perspective, AI-capable AMD desktops do more than improve end-user productivity. The same on-device NPU also supports broader plans for NPU-enabled security, endpoint manageability, and future on-device AI deployments, making the desktop refresh a single investment that serves multiple IT roadmaps at once.

Implications for Office Workers with AI Desktops

At the same time, AI adoption introduces operational considerations that should be addressed during planning rather than after deployment. Local-model inference, where this study has shown gains, keeps data on the device and avoids the cloud egress and licensing costs that come with cloud-only AI workflows. That is a security and cost benefit, but it requires hardware capable enough to deliver acceptable latency. Policies around model selection, model updates, prompt logging, and acceptable-use guidance should be developed in parallel with the hardware rollout. The investment is in the combined hardware-software-policy package, not just the desktops themselves.

Ultimately, the case for AI-capable AMD desktops for office workers is both practical and strategic. The productivity gains shown here already justify the investment, and the rate of improvement across silicon, models, and software ecosystems suggests that future benefits will be substantially larger. Companies that enable AI-enhanced office work now are not just optimizing the efficiency of their largest workforce segment, they are building the foundation for a workforce that can act faster, communicate more clearly, and adapt more effectively to the increasing pace of business in the coming years.

Key Highlights:

The new AMD Ryzen AI 7 PRO 450G desktop **reduces office work time by 66%** in a typical AI-enabled workflow, completing in 19.1 minutes what manual workflows take 56.7 minutes to accomplish.

Individual workflow steps run as much as **17x faster with AI** than manual work, with meeting summarization tasks showing the largest gain.

Office workers using AI tools on a modern AMD desktop **save approximately 1.1 workdays per week** at typical task frequencies, the equivalent of roughly 11 work weeks per year.

New silicon runs the same AI workload **1.8x faster than the prior-generation** AMD desktop, and on certain workflows old silicon makes AI assistance slower than the manual approach.

Appendix: Our Testing Theory

Developing a fair and comprehensive comparison between AI-enabled and traditional workflows requires careful consideration of methodology, metrics, and real-world applicability. Our approach centers on creating reproducible, realistic workflows that mirror actual user behavior across different personas. Rather than constructing artificial benchmarks that might favor one approach over another, we have developed complete task sequences that encompass the full complexity of modern work.

For each scenario we test, we create two distinct implementations: one that leverages currently available AI capabilities and another that follows traditional processes. Both workflows target identical outcomes and deliverables, ensuring that we are measuring different paths to the same destination rather than comparing fundamentally different tasks. This approach acknowledges that AI-enabled workflows may use entirely different applications or involve novel steps that have no traditional equivalent, such as automated content generation followed by human refinement.

In this study we extended our methodology to include a second hardware variable. Each AI-enabled workflow was run on two AMD desktop systems, a current-generation Ryzen AI 7 PRO 450G and a prior-generation Ryzen 7 5700G. The same AI models, the same prompts, and the same input materials were used on both systems. This allows the silicon contribution to be isolated independently of the model and workflow design. The manual workflow stays identical on both systems and provides a stable third reference point against which both AI configurations can be compared.

Our testing philosophy embraces the current reality of hybrid AI deployment. Despite the rapid evolution of AI PCs and the trajectory toward local, on-device processing, today's optimal solutions often combine cloud-based and local compute resources. Four of the five workflows in this study use local large language models running entirely on the desktop. The fifth uses Windows Copilot+ Click-to-Do, which is a Windows feature with its own internal hybrid execution model. We have documented where local versus cloud processing applies for each workflow, allowing readers to make informed decisions about security, privacy, and infrastructure requirements.

Appendix: Our Testing Theory

The metrics we capture extend well beyond time measurements, though efficiency does form the quantitative basis of our analysis. For local-LLM workflows we also captured Time-to-First-Token and generation throughput, which together explain the source of the silicon contribution measured in the workflow timings. These metrics are reported alongside the workflow times in the Hardware Story section.

Our testing infrastructure combines automated scripts with manual execution to ensure both reproducibility and realism. We use AutoHotKey scripts to standardize user interactions, eliminating variability from factors like typing speed or navigation patterns. Python automation handles data collection, timing measurements, and result aggregation. Specialized tools, including Camo Studio for screen recording and playback, LM Studio and AnythingLLM with Lemonade Server for local large language model deployment, and VB Cable for audio routing, ensure comprehensive data capture. Each workflow undergoes a minimum of three iterations, with results averaged to account for system variability and ensure statistical validity.

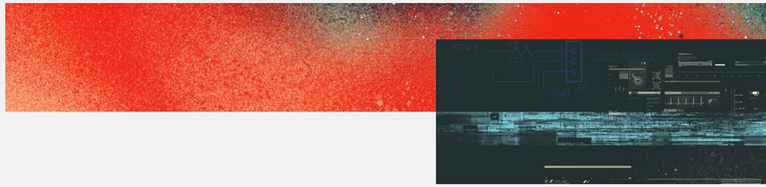

Important Information About this Report

Contact Information
Signal65 | signal65.com | info@signal65.com

Contributors

Ryan Shrout
President & GM - Signal65
Ken Addison
Client Performance Director - Signal65
Ken Fetcher
Performance Analyst and Software Engineer - Signal65

Inquiries

Contact us if you would like to discuss this report and Signal65 will respond promptly.

Citations

This paper can be cited by accredited press and analysts, but must be cited in-context, displaying author's name, author's title, and "Signal65." Non-press and non-analysts must receive prior written permission by Signal65 for any citations.

Licensing

Disclosures

Signal65 provides research, analysis, advising, and lab services to many high-tech companies, including those mentioned in this paper. Research of this document was commissioned by AMD.

In Partnership with:

AMD
together we advance_

About Signal65

Signal65 exists to be a source of data in a world where technology markets and product landscapes create complex and distorted views of product truth. We strive to provide honest and comprehensive feedback and analysis for our clients in order for them to better understand their own competitive positioning and create optimal opportunities to market and message their devices and services.

signal65

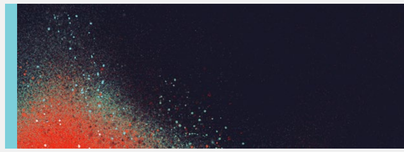

System Configurations

	HP ELITEDESK 8 SFF G2A	HP PAVILION DESKTOP TP01-2XXX
CPU	AMD Ryzen AI 7 PRO 450G	AMD Ryzen 7 5700G
Graphics	AMD Radeon 860M	AMD Radeon Graphics
RAM	32GB DDR5-5600	32GB (16GB x 2) DDR4-2666
Storage	1TB Samsung 980 Pro	1TB Samsung MZVL81T0HELB-00BH1
Display	27" 3840x2160	27" 3840x2160
System BIOS	Y07 92.03.00	F.36
Operating System	Windows 11 26200.8457	Windows 11 26200.8457
Virtualization Based Security	Enabled	Enabled

Applications Used

LMStudio v0.4.12

Microsoft Office 2604

Lemonade Server 10.4.0

AnythingLLM v1.12.1


signal65