Better manage demanding workloads by upgrading to AI PCs powered by AMD Ryzen processors

Using local AI capabilities to accelerate certain tasks lets workers devote more time and thought to solving complex business problems and better handle the stress of endless to-do lists

Software developers and other technical professionals juggle engineering tasks with repetitive administrative duties. AI can greatly accelerate some tasks, freeing time for those that require concentrated thought. When the AI capabilities run locally on their PCs, these workers also benefit from better privacy and data protection, faster response time and offline availability, and even reduced costs.

To build on our findings in a previous study exploring how AI can support technical professionals, we measured the benefits of upgrading to AI PCs powered by AMD Ryzen™ AI processors and using local AI technologies. We performed six test scenarios using three approaches:

- **Using local AI technologies** on a modern AI-capable PC powered by an AMD Ryzen AI processor
- **Completing the same tasks manually** on the same modern AMD Ryzen AI processor–based PC
- **Completing the tasks manually** on a 4-year-old PC

Compared to traditional, non-AI methods on the older PC, using local AI tools on the new AI PC cut task completion times by up to 97.5 percent. In this example, summarizing a meeting transcript took 1.5 minutes instead of an hour. These results illustrate the incredible time savings technology professionals can achieve by using AI-capable PCs powered by AMD Ryzen processors to incorporate AI into their daily workflows.

Tech expert
Performs software development and other tasks

Reclaim up to 2 work days per week

for high-value work by by streamlining activities with local AI features

Cut the time to complete a task

by up to 97.5% by using local AI technologies

The responsibilities a busy developer must manage

Consider a hypothetical professional named Valerie. Despite the fact that she was hired to develop software, and that remains her primary role, it's hard for her to get the uninterrupted time she needs for deep coding work. Like many technical professionals, her schedule can quickly fill up with other necessary activities—such as following industry developments, responding to Jira tickets, and writing documentation.

If she's not careful, these responsibilities can crowd out the focused engineering time required to produce her best work. Delegating some of these routine tasks to AI could help her lower stress, maintain better concentration, and ultimately produce higher-quality results.

The benefits of running AI tools locally on your PC

While many useful AI solutions are delivered through the cloud, running AI applications locally on personal computers can offer several important advantages for workers such as Valerie:

Stronger privacy and data protection. When information remains on the local device, it never needs to travel to external servers. This reduces the risk of data exposure, helps keep proprietary work confidential, and can make it easier to satisfy regulatory or compliance requirements.

Immediate responses and offline availability. Local AI eliminates the need to send requests to remote data centers and wait for responses, which can significantly reduce delays. These tools can also continue functioning without an internet connection, making them dependable in remote locations, during network interruptions, or within secure, air-gapped environments.

Lower long-term costs. Many cloud AI platforms rely on recurring subscription or usage fees. Locally run AI tools generally avoid ongoing service charges, which can make them more economical over time.

How Valerie can work smarter with a new AMD Ryzen AI PC

Compared to their 4-year-old counterparts, new AI PCs equipped with AMD Ryzen AI processors offer considerable improvements. To measure their potential for streamlining tasks and improving life for professionals such as Valerie, we started our testing by carrying out six scenarios on an HP EliteBook X G2a with an AMD Ryzen AI 7 PRO 450 processor:

- Collecting the latest news
- Generating and scheduling tasks from emails
- Summarizing a meeting transcript
- Building a predefined application
- Submitting a new Jira ticket
- Drafting a document

Modern system

HP EliteBook X G2a

AMD Ryzen AI 7 PRO 450 CPU
AMD Radeon™ 860M GPU
32 GB of LPDDR5 memory

We completed each scenario two ways: (1) with local AI tools and (2) following a manual approach. AMD provided initial tasks and workflows that we reviewed and—in some cases—adjusted to reflect real-world usage. For some tasks, a single individual performed the manual approach three times, and for others, three individuals each performed the task once. For complete details on the test systems and our procedure, see the science behind the report.

Table 1 shows our findings. Using AI tools cut as much as 97.5 percent of the time needed to perform these activities. (We summarize the different approaches in the sections to come and provide detailed steps in the science behind the report.)

Table 1: Time to complete tasks on the HP EliteBook X G2a with and without AI to show how AI improves user experience on a modern system. Source: PT.

Task	Time without AI hh:mm:ss	Time with AI hh:mm:ss	Time saved hh:mm:ss	Time saved Percentage
Collecting the latest news	0:50:19	0:22:20	**0:27:59**	**55.6%**
Generating and scheduling tasks from emails	0:07:12	0:00:49	**0:06:23**	**88.6%**
Summarizing a meeting transcript	1:01:04	0:01:31	**0:59:33**	**97.5%**
Building a predefined application	93:20:00	17:02:10	**76:17:50**	**81.7%**
Submitting a new Jira ticket	0:07:06	0:03:10	**0:03:56**	**55.3%**
Drafting a document	0:22:16	0:11:02	**0:11:14**	**50.4%**

How much time could Valerie save by upgrading to a current AI-capable PC powered by an AMD Ryzen processor and using on-PC AI tools?

If Valerie were using a 4-year-old PC, upgrading to a new system could be very beneficial. To explore the potential time savings of such an upgrade, we reran the manual testing using an HP EliteBook 865 16 inch G9 Notebook PC equipped with an AMD Ryzen 5 PRO 6650U processor. In the pages that follow, we compare the time required to complete tasks using AI on the current system with the time it takes using a manual approach on the older system.

2026 system

HP EliteBook X G2a[1]

AMD Ryzen AI 7 PRO 450 CPU
AMD Radeon 860M GPU
32 GB of LPDDR5 memory

2022 system

HP EliteBook 865 16 inch G9 Notebook PC

AMD Ryzen 5 PRO 6650U CPU
AMD Radeon GPU[2]
32 GB of DDR5 memory

About the AMD Ryzen processors in our test systems

The HP EliteBook X G2a we tested featured an AMD Ryzen AI 7 PRO 450 processor with AMD Radeon 860M GPU. It had 8 cores, 16 threads, a 5.1GHz Max boost clock, 8MB L2 cache, 16MB L3 cache, 32 GB of DDR5 memory, and up to 50 trillions of operations per second (TOPS) from the NPU and up to 66 TOPS total AI performance.

▸ Learn more

The HP EliteBook 865 16 inch G9 Notebook PC we tested featured an AMD Ryzen 5 PRO 6650U processor. It had 6 cores, 12 threads, a 4.5GHz Max boost clock, 3MB L2 cache, 16MB L3 cache, and 32 GB of DDR5 memory.

▸ Learn more

Collecting the latest news

Part of Valerie's responsibilities is staying informed about industry news and ensuring that her team remains current as well. For the AI-driven version of this scenario, which we performed on the current AI PC powered by an AMD Ryzen AI processor, we used the Generate feature with the Qwen3 4B model to import curated news articles on a specific topic (AMD in our test), create summaries for each article, and assemble those summaries into a Microsoft Word document for sharing with colleagues.

For the manual version of this workflow, which we conducted on the 4-year-old system, the process involved searching online for recent AMD-related news, opening five articles, reading through them, taking notes, and then writing a summary in Word by hand. Three participants carried out the manual process, and we report the average time they required to complete the task.

As Figure 1 shows, **using the AI feature cut the time to complete this task in half,** saving an average of 56.2 percent of the work time—over 28 minutes of time savings for this task alone.

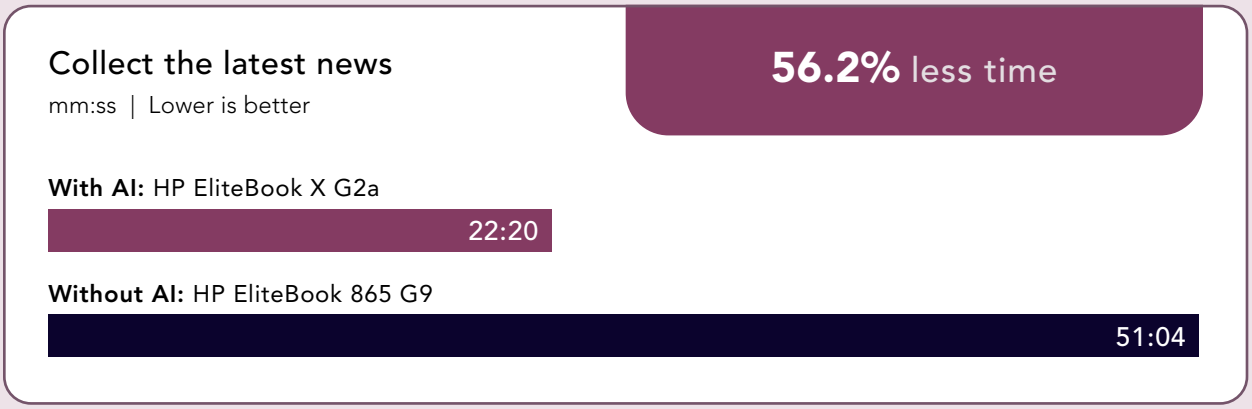

Figure 1: Time to collect news on a current AMD Ryzen AI processor-powered PC with AI assistance and on a 4-year-old PC without AI assistance. Source: PT.

Note that the graphs in this report use different scales to keep a consistent size. Please be mindful of each graph's data range as you compare.

Generating and scheduling tasks from emails

Valerie's manager often assigns work through email messages. For the AI-enabled version of this scenario, which we performed on the current AI PC powered by an AMD Ryzen AI processor, we used Noodle, an AI-powered productivity tool that manages work, email, and calendars in a single, unified platform. (Note: We used an alpha version that AMD provided.) We used Noodle with a local Llama model to analyze the email, identify the required subtasks, automatically generate a task, and convert it into a scheduled calendar event.

For the manual version of the workflow, which we conducted on the 4-year-old system, the process involved reading the email in Microsoft Outlook, determining the required tasks manually, creating those tasks by hand, and scheduling them directly in the Outlook calendar. We repeated the task three times for each method and report the average completion time.

Using the AI feature greatly shortened the time required to complete the task, **reducing a 7-minute task to less than a minute**—an improvement of 88.8 percent (see Figure 2).

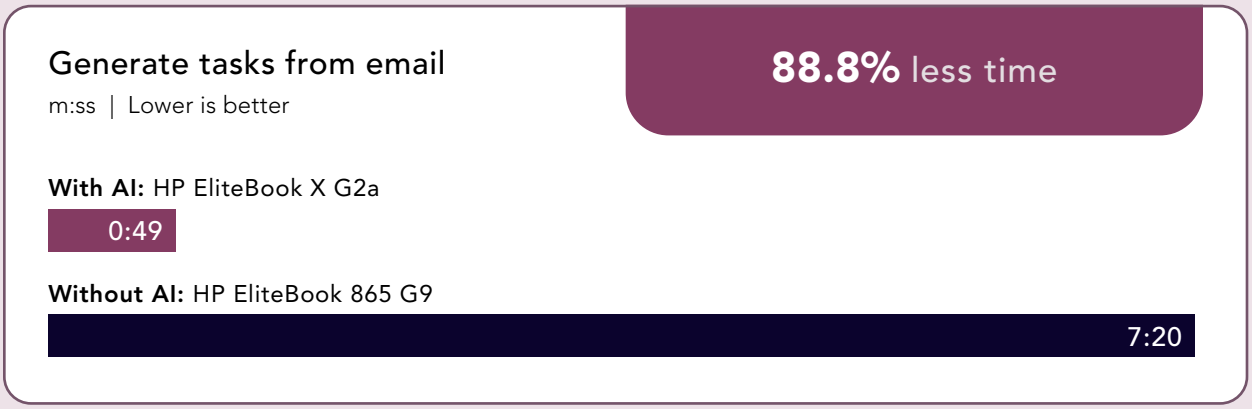

Figure 2: Time to generate tasks from emails on a current AMD Ryzen AI processor-powered PC with AI assistance and on a 4-year-old PC without AI assistance. Source: PT.

Summarizing a meeting transcript

Valerie occasionally receives a meeting transcript, either because she missed the meeting or because a client has provided it as background on their wishes for a project. We tested two ways of summarizing the content of a transcript. When performing this task on the current AMD Ryzen AI processor-powered PC with AI, we used Generate, a secure, local AI platform by Iterate.ai that lets companies run large language models (LLMs) and retrieval-augmented generation (RAG) locally on AI PCs or private servers.

Our scenario used Generate with a Liquid transcript model to ingest the meeting notes file and automatically produce structured meeting minutes. The manual workflow, which three people carried out using the 4-year-old PC, involved reading the meeting transcript, identifying key sections, and drafting the meeting minutes manually in Word using the required structure.

As Figure 3 shows, using the Generate and Liquid AI tools created the notes in **a minute and a half vs. more than an hour** with the manual approach—a savings of 97.5 percent.

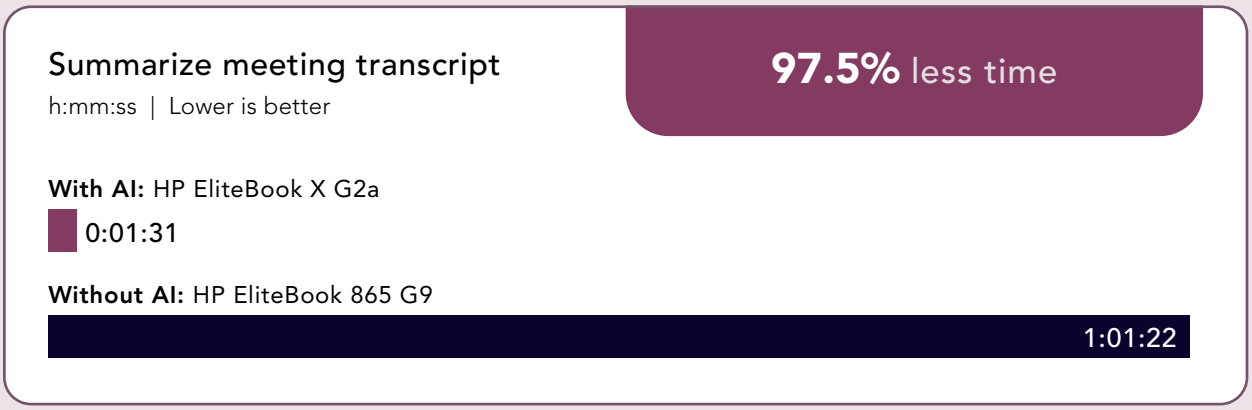

Figure 3: Time to summarize a meeting transcript on a current AMD Ryzen AI processor-powered PC with AI assistance and on a 4-year-old PC without AI assistance. Source: PT.

Building a predefined application

AI-assisted "vibe coding" tools use natural language prompts to produce working code. These applications allow users to create software in a more intuitive way and can significantly reduce development time. For the AI-enabled version of this scenario, which we conducted on the current AI PC powered by an AMD Ryzen AI processor, we used GitHub Copilot together with a Lemonade-hosted large language model to generate and refine the HTTP server code throughout the development process. We repeated the task three times and report the average completion time.

For the manual version of this scenario on the 4-year-old PC, we asked three experienced developers to estimate how long it would take them to build the same predefined application without AI assistance. The assignment required researching, designing, and implementing a JavaScript HTTP server that included basic routing, support for GET and POST requests, JSON response handling, CORS headers, error and 404 responses, and graceful shutdown behavior. We report the average of their estimates.

As Figure 4 shows, our experts estimated that performing the task without AI would take an average of 93 hours—**more than two 40-hour workweeks vs. only 17 hours** with AI. That's a savings of 81.7 percent.

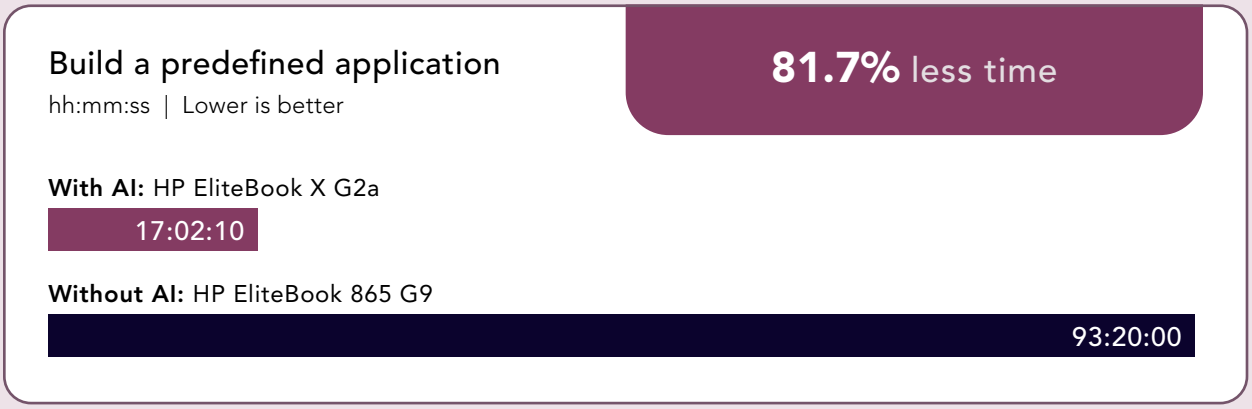

Figure 4: Time to build a predefined application on a current AMD Ryzen AI processor-powered PC with AI assistance and on a 4-year-old PC without AI assistance. Source: PT.

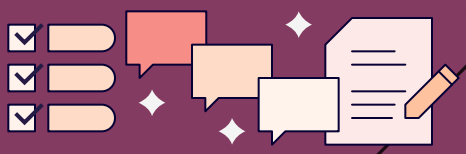

Submitting a new Jira ticket

Valerie's organization relies on Jira, Atlassian's issue tracking and project management platform, to organize and manage work throughout the software development lifecycle. Jira tickets help teams document, track, and resolve tasks as projects move forward. This test scenario compares two different methods of creating Jira tickets.

For the AI-driven workflow, which we conducted on the current AI PC powered by an AMD Ryzen AI processor, we used JIRA MCP within Generate AI PC along with a Lemonade-hosted model to produce a fully completed Jira ticket from a single natural-language prompt. On the 4-year-old system, we followed a manual process that involved opening Jira, entering information into each field individually, writing the complete issue description, reviewing the details, and submitting the ticket manually. Three participants carried out the manual process, and we report the average of their completion times.

Using the AI feature reduced the time needed to finish this task by 55.2 percent, **cutting a 7-minute process in half to just over 3 minutes** (see Figure 5).

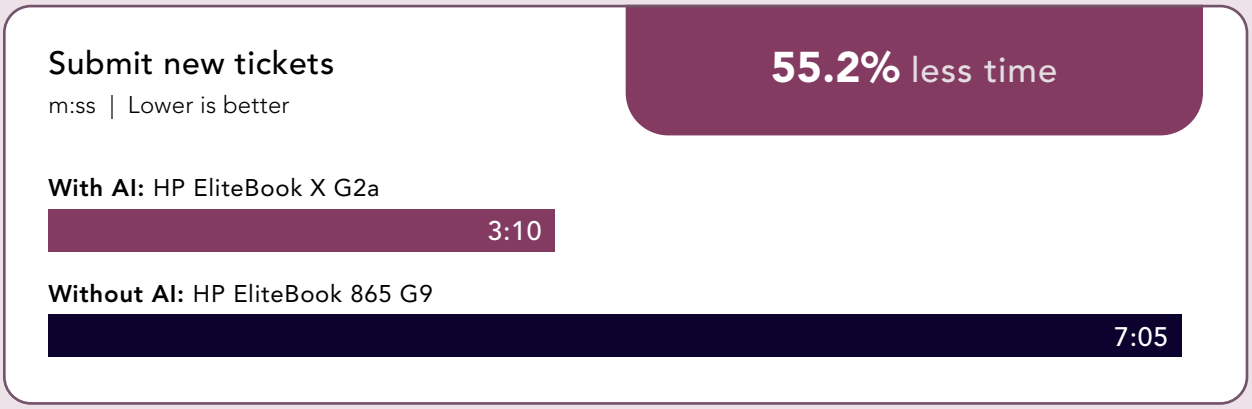

Figure 5: Time to submit a new Jira ticket on a current AMD Ryzen AI processor-powered PC with AI assistance and on a 4-year-old PC without AI assistance. Source: PT.

Drafting a document

As part of her development responsibilities, Valerie frequently produces written materials such as release notes, changelogs, and client-facing project updates. In this scenario, we examined the difference between an AI-assisted writing process and a traditional manual drafting workflow. Although the specific documents professionals create may vary, the overall document-creation process is often similar across roles.

For the AI-assisted workflow on the current AI PC powered by an AMD Ryzen AI processor, we used LM Studio with a Gemma-based model to produce a full press release draft from a prompt and an attached product specification document. In the manual workflow on the 4-year-old system, the user opened Word, reviewed the product specification document, and wrote the entire press release draft manually, including placeholders and supporting details.

Using the AI tools reduced the time required for the task by more than half: **12 minutes vs. 23 minutes** for the manual process, a savings of 52.3 percent.

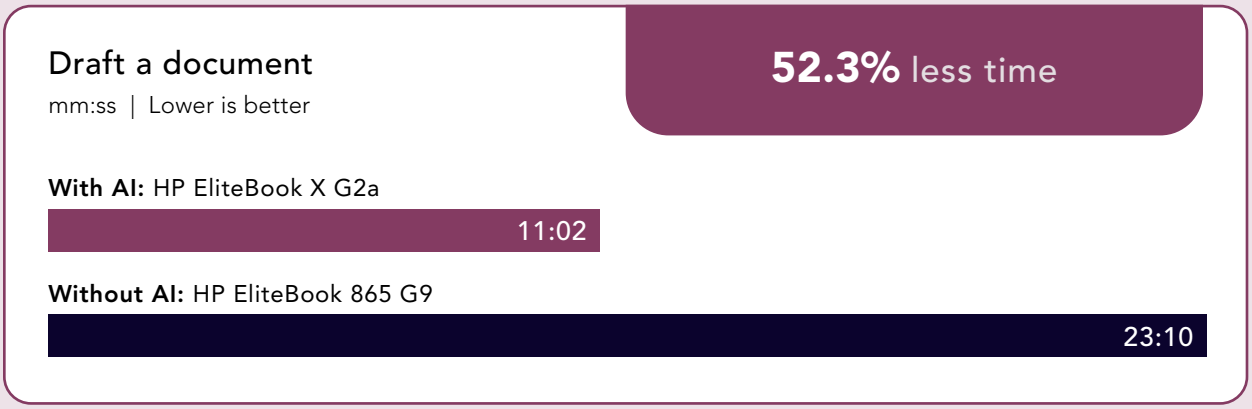

Figure 6: Time to draft a document on a current AMD Ryzen AI processor-powered PC with AI assistance and on a 4-year-old PC without AI assistance. Source: PT.

Time savings across the workflow

Figure 7 shows the total time it took to complete each scenario once. Using AI technologies cut the time dramatically, by 81.5 percent. Obviously, the time savings on the application development activity accounts for the vast majority of this time, just as it would in the workweek of a developer like Valerie.

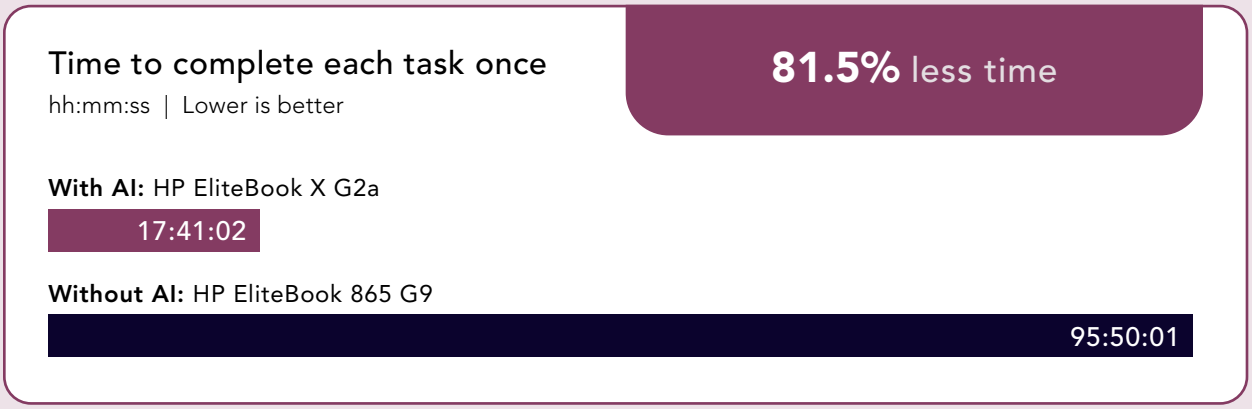

Figure 7: Estimated time to complete all six activities one time on a current AMD Ryzen AI processor-powered PC with AI assistance and on a 4-year-old PC without AI assistance. Source: PT.

How much time would Valerie save over a workweek?

To quantify the time savings a tech expert such as Valerie could experience in a typical workweek, we estimated how many times per week she is likely to execute the tasks we investigated. In our extrapolation, we define a month as four weeks. Based on our developer team's real-life experience, we assumed the following frequencies:

- **Collecting the latest news** ..... 1x daily
- **Generating and scheduling tasks from emails** ..... 5x daily
- **Summarizing a meeting transcript** .. 1x weekly
- **Building a predefined application** ..... 0.5x monthly
- **Submitting a new Jira ticket** ..... 5x daily
- **Drafting a document** ..... 3x weekly

Figure 8 shows the results of our estimations. Based on the frequencies we note above and the time that each approach required, employing AI features on the systems we tested could mean that Valerie would spend just over one-quarter of the time on these routine activities. **Instead of three 8-hour workdays, she could spend only 6 hours and 13 minutes.** This savings of 74.1 percent would allow her to take on more high-value work and solve higher-level problems.

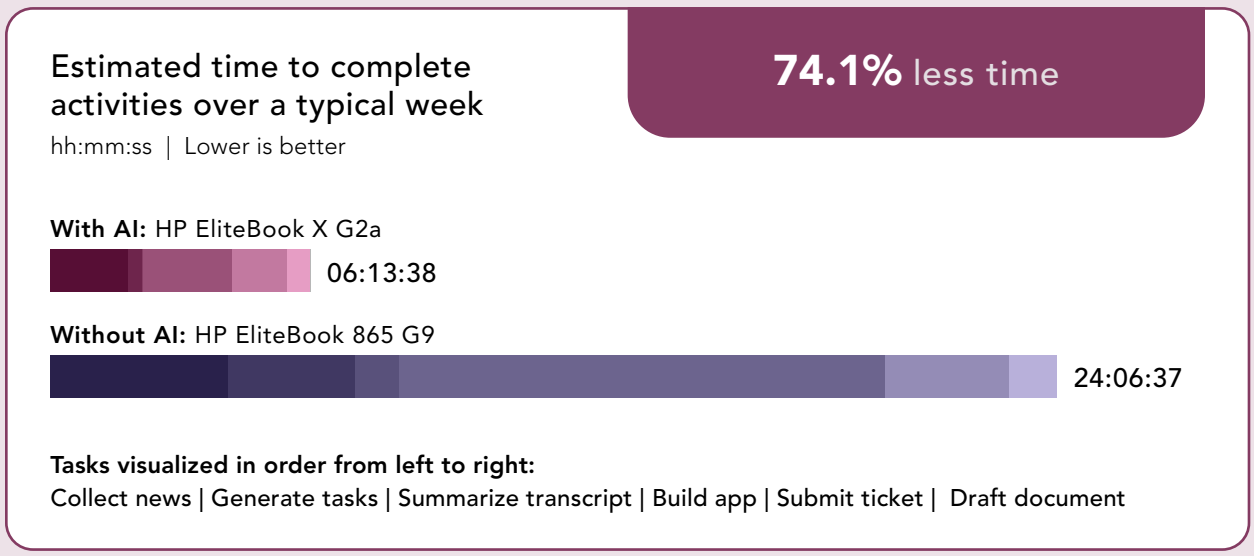

Figure 8: Estimated time to complete all six activities over a week on a current AMD Ryzen AI processor-powered PC with AI assistance and on a 4-year-old PC without AI assistance. Source: PT.

Conclusion

Using AI capabilities on modern PCs can greatly improve efficiency for technical professionals, letting them shift the time they would have spent on certain routine tasks to higher-value activities that demand uninterrupted focus. By choosing AI tools that run locally rather than in the cloud, these workers can keep proprietary data secure, enjoy faster response time and work offline if necessary, and avoid recurring subscription or usage fees.

In our testing, using local AI tools on a new HP EliteBook X G2a AI PC powered by an AMD Ryzen AI processor reduced task completion times by as much as 97.5 percent compared to completing the same tasks manually on a 4-year-old HP EliteBook 865 16 inch G9 Notebook PC.

Based on our estimates of how frequently these tasks typically occur, we found that using AI tools on the newer system could lower weekly development-related work time from an average of 24 hours and 6 minutes to approximately 6 hours and 13 minutes. This represents a **reduction of almost 18 hours of work per week—equivalent to over two standard workdays**.

Our testing also showed that AMD Ryzen AI PRO 400 Series processors effectively supported enterprise AI workloads. With the AI capabilities on these new PCs, workers like Valerie could use this reclaimed time to concentrate on higher-value responsibilities.

1. The HP EliteBook X G2a we tested was a pre-release model. It is possible that it will look different from the model pictured.
2. The HP EliteBook 865 G9 did not report the specific model of the GPU.

This project was commissioned by AMD.

Read the science behind the report ▶

Primary contributors

Tech: Jesse R., Travis B.
Writing: Laura W.
Design: Jared White
PM: Scott Luchene

How we created this report

A PT team, which includes the contributors we've listed and others, created this report and performed the technical work behind it. In addition to our use of AI in testing, we used AI to draft sections of this report.

Principled Technologies®

Facts matter.®