



# Greatly accelerate complex tasks by upgrading to AI PCs featuring AMD Ryzen processors

Using AI features on an AMD Ryzen processor-powered PC let us execute coding and other tasks in substantially less time than using conventional, manual methods on a 3-year-old system

Many software developers and other technical professionals split their days between focused engineering work and routine tasks, such as meetings and email. By applying local AI technologies, they can spend less time on both coding and administrative activities. This lets them devote more time and mental bandwidth to deeper problem solving and can reduce the stress caused by an ever-growing to-do list.

As a follow-up to a [previous study](#) investigating the ways that AI can benefit technical professionals, we conducted tests to quantify the benefits of upgrading to new AI PCs featuring AMD Ryzen™ processors. We ran a set of defined scenarios three ways:

- **Using AI technologies** on a current AI-capable PC powered by an AMD Ryzen processor
- **Using a manual approach** on a current AI-capable PC powered by an AMD Ryzen processor
- **Using a manual approach** on a 3-year-old PC

Across the tasks we tested, using local AI features reduced task completion time by as much as 97.6 percent compared with traditional, non-AI approaches. Our results indicate AI-capable PCs powered by AMD Ryzen processors are a strong option for technology professionals who want to take advantage of AI in their daily workflows.



**Tech expert**  
*Performs software development and other tasks*

**Reclaim up to 2 workdays per week for high-value work**

*by employing local AI features to help streamline activities*

**Cut the time to complete a task**

*by up to 97.6% by using AI technologies*

## The tasks that a busy software developer must juggle



Let's imagine a hypothetical professional named Ravi. While his official title is software developer, he sometimes struggles to achieve the focused blocks of coding time he needs to do his best work. If he's not careful, the other necessary tasks that pop up regularly—things such as keeping up on industry news, dealing with Jira tickets, and drafting documentation—can fill his days. Offloading some of this to AI could reduce his stress levels, improve his focus, and help him deliver higher-quality work.

## The advantages of local AI tools

While many cloud-based AI tools are valuable, by choosing to use AI tools that run locally on your PC, you reap a number of benefits.

**Better privacy and data security.** Data never leaves your machine, eliminating risks of data breaches; your confidential, proprietary work stays that way; and it's easier to meet regulatory compliance.

**Zero-latency and offline functionality.** Eliminating back-and-forth communication with a distant data center, local AI can provide near-instant responses. Local tools also work without an internet connection, making them reliable for remote locations, during outages, or in secure, air-gapped environments.

**Long-term cost efficiency.** In contrast to ongoing, expensive monthly subscriptions with many cloud-based AI services, local AI has no recurring usage costs.

## How a new AMD Ryzen AI PC can help Ravi work smarter

New AI PCs equipped with AMD Ryzen processors offer considerable improvements over their 3-year-old counterparts. To understand how they streamline tasks and improve life for professionals such as Ravi, we started our testing by carrying out six scenarios on an ASUS™ Zenbook™ S16 powered by an AMD Ryzen AI 9 HX 470 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



**Current system**  
**ASUS Zenbook S16**  
 AMD Ryzen AI 9 HX 470 processor  
 AMD Radeon 890M GPU  
 32 GB of DDR5 memory

We completed each scenario two ways, using local AI tools and using 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 compares using AI tools and a manual method on a new AI PC with an AMD Ryzen processor. Using AI tools cut as much as 97.6 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 ASUS Zenbook S16 with and without AI to show how AI improves user experience on a current 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	0:50:35	0:23:39	<b>0:26:56</b>	<b>53.2%</b>
Generating and scheduling tasks from emails	0:08:50	0:01:06	<b>0:07:44</b>	<b>87.5%</b>
Summarizing a meeting transcript	1:01:18	0:01:27	<b>0:59:51</b>	<b>97.6%</b>
Building a predefined application	93:20:00	17:03:01	<b>76:16:59</b>	<b>81.7%</b>
Submitting a new Jira ticket	0:07:06	0:03:11	<b>0:03:55</b>	<b>55.1%</b>
Drafting a document	0:22:27	0:11:38	<b>0:10:49</b>	<b>48.1%</b>

## What kind of time savings could Ravi achieve by upgrading to a current AI-capable PC powered by an AMD Ryzen processor and using local AI tools?

Imagine that Ravi is using a 3-year-old PC. To explore that, we re-ran the manual approach testing using an ASUS Zenbook 15 powered by an older AMD Ryzen processor. In the sections that follow, we quantify the time savings he could enjoy by comparing the time to complete tasks using AI on the current system vs. using a manual approach on the 3-year-old system.



Current system

### ASUS Zenbook S16

AMD Ryzen AI 9 HX 470 processor  
AMD Radeon 890M GPU  
32 GB of DDR5 memory



3-year-old system

### ASUS Zenbook 15

AMD Ryzen 7 7735U processor  
AMD Radeon 680M GPU  
32 GB of DDR5 memory

## About the AMD Ryzen processors in our test systems

The **ASUS Zenbook S16 we tested** featured an AMD Ryzen AI 9 HX 470 processor. It had 12 cores, 24 threads, a 5.2 GHz boost clock, a 12MB L2 cache, a 24MB L3 cache, 32 GB of DDR5 memory, and 86 trillions of operations per second (TOPS) total AI performance.

► [Learn more](#)

The **ASUS Zenbook 15 we tested** featured an AMD Ryzen 7 7735U processor. It had 8 cores, 16 threads, a 4.75 GHz boost clock, a 4MB L2 cache, a 16MB L3 cache, and 32 GB of DDR5 memory.

► [Learn more](#)



## Collecting the latest news

Part of Ravi's role is keeping himself and his team up to date on industry developments. For the AI execution of this scenario, which we completed on the current AMD Ryzen processor-powered ASUS Zenbook S16, we used Generate with the Qwen3 4B model to load curated news articles related to a specific topic (in our case, AMD), generate summaries for each, and compile them into a Word document for distribution to colleagues. We carried out the manual workflow for this scenario on the 3-year-old ASUS Zenbook 15. That required searching the web for recent AMD news, opening five articles, reading them, taking notes, and manually writing a summary in Word. Three individuals completed the activity manually, and we report the average of their times.

As Figure 1 shows, using the AI feature **cut the time to complete this task in half**, saving an average of 53.6 percent of the work time—over 27 minutes of time savings for this task alone.

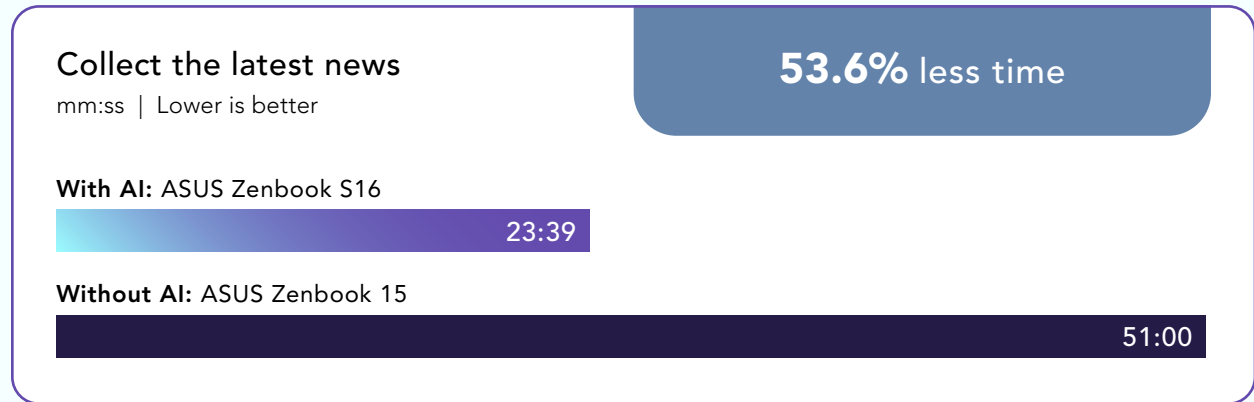



Figure 1: Time to collect news on a current AMD Ryzen processor-powered PC with AI assistance and on a 3-year-old PC without AI assistance. Source: PT.

Please 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

Ravi's manager frequently uses emails to assign tasks. When we carried out the AI workflow for this scenario on the current AMD Ryzen processor-powered PC, we used Belt with a local Llama model to break down the email into subtasks, auto-generate a task, and convert it into a scheduled event. We again used the ASUS Zenbook 15 for the manual workflow, which required reading the email, identifying tasks manually, creating tasks by hand, and scheduling them directly in the Microsoft Outlook calendar. We performed the task three times using each approach and report the average time.

Using the AI feature cut the time to complete this scenario from **almost 8 minutes to just over 1 minute**, an average of 86.0 percent (see Figure 2).

### Generate tasks from email

m:ss | Lower is better

**86.0%** less time

**With AI:** ASUS Zenbook S16

1:06

**Without AI:** ASUS Zenbook 15

7:52

Figure 2: Time to generate tasks from emails on a current AMD Ryzen processor-powered PC with AI assistance and on a 3-year-old PC without AI assistance. Source: PT.



## Summarizing a meeting transcript

Ravi occasionally receives a meeting transcript, either because he 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 processor-powered PC with AI, we used Generate AIPC 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 ASUS Zenbook 15, involved reading the meeting transcript, identifying key sections, and drafting the meeting minutes manually in Word using the required structure. We use the average of these times.

As Figure 3 shows, using the Generate and Liquid AI tools to create the notes took **a minute and a half vs. the 1 hour** necessary with the manual approach.

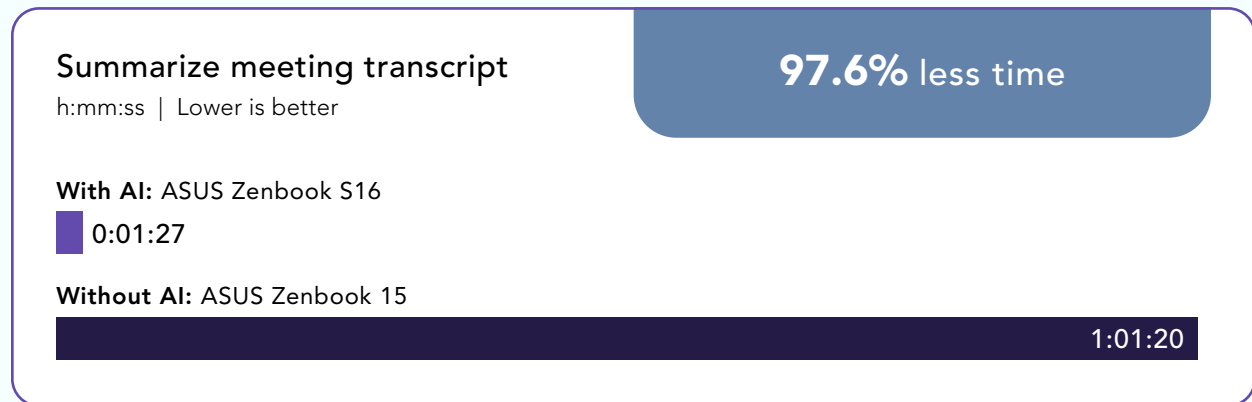


Figure 3: Time to create meeting notes on a current AMD Ryzen processor-powered PC with AI assistance and on a 3-year-old PC without AI assistance. Source: PT.



## Building a predefined application

Vibe coding applications are AI-powered tools that use natural language prompts to generate code. They enable users to build software intuitively and have the potential to be great time savers. In the AI approach to this scenario on the current AMD Ryzen processor-powered PC, we used GitHub Copilot assisted by a Lemonade-hosted LLM to generate and refine the HTTP server code during the development process. We performed the task three times and report the average time.

For the manual approach on the 3-year-old PC, we asked three experienced developers to estimate the amount of time it would take them to develop a predefined application without AI assistance. The task involved researching, designing, and implementing a JavaScript HTTP server with basic routing, GET/POST support, JSON response handling, CORS headers, error/404 responses, and graceful shutdown. We report the average of their estimates.

As Figure 4 shows, our experts estimated that the task would take an average of 93 hours. This is the equivalent of **more than 2 weeks of 8-hour workdays vs. less than half of a typical workweek** with the AI approach.

### Building a predefined application

hh:mm:ss | Lower is better

81.7% less time

**With AI:** ASUS Zenbook S16

17:03:01

**Without AI:** ASUS Zenbook 15

93:20:00

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



## Submitting a new Jira ticket

Ravi's company uses Jira, an issue tracking and project management tool from Atlassian. Jira tickets organize, track, and manage tasks throughout the project lifecycle. This test scenario compares two ways of submitting Jira tickets. The AI workflow, which we carried out on the current AMD Ryzen processor-powered PC, used JIRA MCP inside Generate AIPC with a Lemonade-hosted model to create a fully populated Jira Software ticket from a single natural-language prompt. On the 3-year-old system, we executed the manual workflow: opening Jira, filling out all fields manually, typing the full issue description, reviewing the details, and submitting the ticket by hand. We performed the task three times using each approach and report the average time.

Employing the AI feature **cut the time required to complete this task by more than half**, saving 4 minutes (see Figure 5).

### Submit new tickets

m:ss | Lower is better

55.6% less time

**With AI:** ASUS Zenbook S16

3:11

**Without AI:** ASUS Zenbook 15

7:11

Figure 5: Time to submit a new Jira ticket on a current AMD Ryzen processor-powered PC with AI assistance and on a 3-year-old PC without AI assistance. Source: PT.



## Drafting a document

For his development work, Ravi often writes documentation such as release notes, changelogs, and project updates for clients. This test scenario explored the difference between an AI-enhanced composition workflow and the traditional manual drafting process. Every professional writes different types of documents, but the document creation workflow can be similar across roles. The AI workflow on the current AMD Ryzen processor-powered PC used LM Studio with a Gemma-based model to generate a complete press release draft from a prompt and attached product specification document. The manual workflow on the ASUS Zenbook 15 required opening Word, referencing the product specification document, and manually writing the full press release draft with placeholders and details. We performed the task three times using each approach and report the average time.

Using the AI tools **cut 20 minutes off a 32-minute task** when done manually, a savings of 63.6 percent.

### Draft a document

mm:ss | Lower is better

**63.6% less time**

**With AI:** ASUS Zenbook S16

11:38

**Without AI:** ASUS Zenbook 15

32:02

Figure 6: Time to draft a document on a current AMD Ryzen processor-powered PC with AI assistance and on a 3-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 from **more than 2 workweeks to less than 18 hours**, an average of 81.5 percent. Obviously, the time savings on the application development activity accounts for the vast majority of this time, as would be the case in the workweek of a developer like Ravi.

### Time to complete each task once

hh:mm:ss | Lower is better

**81.5%** less time

**With AI:** ASUS Zenbook S16

17:44:02

**Without AI:** ASUS Zenbook 15

95:59:25

Figure 7: Estimated time to complete all six activities one time on a current AMD Ryzen processor-powered PC with AI assistance and on a 3-year-old PC without AI assistance. Source: PT.



## Potential time savings across a workweek

To get a sense of the kind of time savings that a tech expert could experience over the course of a typical workweek, we estimated how many times per week they might execute the tasks we investigated. For the purposes of this extrapolation, we define a month as four weeks. We assumed the following frequencies based on our developer team's real-life experience:

- **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 Ravi would spend approximately one-quarter of the time on these routine activities: **3 typical 8-hour workdays vs. less than a single workday**. This would allow him to take on more high-value work and solve higher-level problems.

### Estimated time to complete activities over a typical week

hh:mm:ss | Lower is better

**73.8%** less time

**With AI:** ASUS Zenbook S16



**Without AI:** ASUS Zenbook 15



**Tasks visualized in order from left to right:**

Collect news | Generate tasks | Summarize transcript | Build app | Submit ticket | Draft document

Figure 8: Estimated time to complete all six activities over a week on a current AMD Ryzen processor-powered PC with AI assistance and on a 3-year-old PC without AI assistance. Source: PT.

## Conclusion



Technical professions can leverage AI capabilities on modern PCs to improve reallocate time they would have spent on certain routine tasks to higher-value activities that require uninterrupted focus. By selecting 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.


We investigated the potential time savings for software developers and other tech users by utilizing AI technologies on AI-powered PCs with AMD Ryzen processors. In our testing, using local AI technologies on a new ASUS Zenbook S16 system reduced task completion times by as much as 97 percent compared to using manual processes on a 3-year-old ASUS Zenbook 15. Based on our estimations of the frequency with which the test tasks would occur, we determined that using AI technologies on the new PC could reduce the weekly development work time from an average of 24 hours and 48 minutes to 6 hours and 29 minutes on average—a **savings of over 18 work hours, or more than two typical workdays.**


Workers like Ravi could repurpose this time saved to focus on higher-value activities. As our testing showed, AMD Ryzen AI PRO 400 Series processors did a strong job of supporting enterprise AI workloads.


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:** Trent Weatherford

### 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.



**Facts matter.®**

Principled Technologies is a registered trademark of Principled Technologies, Inc. All other product names are the trademarks of their respective owners. For additional information, review the science behind this report.