


SAPPHIRE EDGE AI MINI PC



www.sapphiretech.com

THE TIME FOR AI IS NOW

The technology industry is moving at an increasingly frenetic pace driven by the relentless progress of artificial intelligence (AI). Though the aim is for AI to be democratised for everyone, available anywhere and at any time, only a few companies currently dictate its evolution.

On the one hand, you have producers of the hardware required to run at-scale AI training and inference, where the likes of NVIDIA and AMD are making bank to the tune of billions of dollars. Then there are companies optimising large-language models (LLMs) that run on said hardware, so think of OpenAI, Anthropic, Meta, Microsoft, et al, whose attendant models like ChatGPT and Copilot+ are known to everyone.



Yet the use of AI itself is undergoing a transformation as existential issues such as data security and integrity become increasingly important, especially when these

multi-billion-parameter LLMs are hosted in the cloud and physically resident in data centres worldwide. Put simply, the data you're relying on for AI is oftentimes located far, far away.

This is precisely why having an AI-capable PC, running models locally and sandboxing your sensitive data in a secure environment, is gaining traction amongst personal users and businesses alike.

The question is simple enough: would you prefer to have the nuts and bolts driving creative AI endeavours within your control, on your device, or beamed across the Internet? As a business, you wouldn't want sensitive financial data, from which, say, you build company-specific AI forecasts, anywhere but on your local hard drive.

Running local LLMs has other advantages. Reliability is enhanced because you don't need an Internet connection to generate answers. Everything is done locally, meaning responses are predictable and not at the whims of cloud providers who may queue up work.




AI PCs are built different

Serving this need for enhanced privacy is the domain of nascent AI PCs. At their heart, they are regular computers that you know and love - be they laptops or desktop machines - but carry specific hardware primed for running the increasingly complicated maths behind using AI at acceptable speeds. In effect, this hardware, typically known as a neural processing unit (NPU), does the heavy AI lifting that almost magically builds content for your marketing department, creates arresting videos, or helps in tidying up emails and presentations, to name but a few scenarios. The applications for AI in the workplace or home are almost endless, and you need somewhere to run it.

To be clear, local AI infrastructure is fundamentally less complex than its cloud-based counterparts, but that's perfectly okay as AI PCs are endpoints. They don't need extensive hardware for the computationally-heavy part of AI known as training. Rather, functioning as inference engines that use pre-trained data contained within downloadable models, they are suitably powerful to offer a seamless real-time experience. If you want to think of it in a more readily accessible way, the training part of AI is like learning how to play chess; it takes thousands of hours to become good. Inference, on the other hand, is like playing an actual game using the skills you have developed. It's easier to do and a natural consequence of the other.

Small size, big potential

Running local AI has never been easier, and there's absolutely scope to do it on a mini PC that's secure, convenient and powerful.



HUMAN-LIKE INTERACTION

Thinking more deeply about AI PCs, the combination of a resident NPU - supported by a capable CPU and GPU - and a decent amount of onboard system memory to house AI models is enough for a fluid experience. Running on such hardware, ask a well-tuned local AI model a question and it will respond within a second or two, known in the industry as time to first token. Once a system has achieved a good token throughput rate, it feels more like a flowing conversation than a stilted back-and-forth that plagued early iterations of AI.

Recent evolution of these models is startling. Just a couple of years' ago, an AI conversation felt like talking to a child who's having a bad day. Now, AI models that work well on a local machine exhibit a level of pseudo intelligence and understanding that's akin to chatting with an undergraduate student. As fine-tuning and novel ways are found to elevate the experience further, the sky is the limit.

The purpose of on-device AI is to make your life simpler and more engaging for both work and play whilst maintaining robust security and low latency, all without ever having to use the cloud in a meaningful way. Generally speaking, the greater the number of parameters contained in a certain model, the better it will be at tasks such as summarising emails and documents, offering writing suggestions, real-time text generation (the two-way conversation alluded to above), translating languages, and building marketing materials through parsing heaps of relevant data. You know, the general AI goodness we're quickly becoming accustomed to.



One of the most visible examples of local AI is Microsoft's Copilot+, and it carries similar functionality as massive cloud-based solutions like the oft-mentioned ChatGPT. Microsoft knows a thing or two about bringing AI to life and is acutely aware that first experiences matter hugely. There's little point in local AI if it takes many minutes to find a particular email or answer a taxing question. This is why Copilot+ certification demands your AI PC have an NPU with a capability of at least 40 trillion operations per second (TOPS) and system memory of at least 16GB.

It's all well and good explaining how productivity-enhancing AI can run well on your PC, but the real proof of the pudding, as they say, is in the eating. To that end, the versatility of an AI PC needs to be demonstrated, reinforcing its usefulness to a much broader audience.

THE ANSWER

As a case study, it's instructive to examine the latest Sapphire Edge AI mini PC series, which promises the power necessary for the modern era in a conveniently tiny package. Available in a bare-bones configuration, the key attribute that makes it AI-ready is the potent processor contained within.

Each Edge AI mini PC carries an NPU capable of performing up to 50 TOPS alone, with a further 30 or so TOPS available from the CPU and GPU, depending upon configuration.

All models have the same NPU capability - the driving force behind processing AI locally - but differ in the specification of the AMD Ryzen processor. Starting at just £439 for the base model, where you merely add memory and storage of your choice, you don't need to spend a fortune in building an AI-first PC that's small yet eminently powerful.

AI Made Simple

This is a handsomely petite PC, measuring 117 x 111 x 30mm and weighing just 415g. Suffice to say, this system slots into practically any space without much of a footprint.

You can also tuck it away behind a monitor via VESA mounting if you'd prefer. It's easy to see why the mini PC boom is truly underway; for many an office, these space-saving, low-power and high-performance credentials make for an attractive replacement for a traditional desktop tower.



Think of it as the guts of a laptop in a tiny little box that's particularly adept at AI. Slipping into a bag without any fuss, it's a portable solution for millions who hybrid work between the home and office.

The sides feature small but plentiful ventilation cutouts amidst golden grills. Look a little closer, though, and you'll find a Kensington Security Slot. The black-gold colour scheme Sapphire adopts for its Edge AI PCs is classy, too.

Accessing the majority of Edge AI mini PC's internals is toolless and simple. The magnetic cover is easy to remove, and once it's off both RAM and storage slots reveal themselves for effortless deployment at scale.

You can go from a barebone unit to fully assembled, AI-capable PC in the time it takes to boil a kettle.

On-device Smarts

Local AI may have felt like a fad a short while back, but now it has meaningful benefit for businesses who must innovate or be left behind.

There's absolutely every reason why your next personal or business PC needs to have AI smarts baked right in. Sapphire Edge is a capable system first and foremost, but AI hardware turbo-charges its utility right through from one-man-band businesses to hundreds of units for large corporations.

Modern Specifications

The Sapphire Edge AI mini PC is a compact system with a total volume of under 0.4 litres. That's less than a can of beer, which isn't designed for AI-driven tasks, content creation, productivity, and edge computing.

It features AMD Ryzen AI 300 Series processors with three variants: the flagship Edge AI 370 powered by the 12-core Ryzen AI 9 HX 370 with Radeon 890M graphics, the mid-range Edge AI 350 using the 8-core Ryzen AI 7 350 with Radeon 860M, and the entry-level Edge AI 340 equipped with the 6-core Ryzen AI 5 340 and Radeon 840M. All models include a dedicated XDNA2 NPU delivering up to 50 TOPS of AI performance, enabling efficient local inference and Copilot+ PC compatibility.


The systems support up to 96GB of DDR5 SO-DIMM memory via dual slots and offer expandable storage through two M.2 PCIe NVMe slots (one 2280 and one 2242), plus an additional M.2 2230 slot for Wi-Fi 6E and Bluetooth. Connectivity is robust with dual USB-C 4.0, three USB 3.2, a single USB 2.0, plus HDMI 2.1, 2.5G Ethernet, and a 3.5mm audio jack.



WHY I PREFER LOCAL AI

It's incredibly easy to get up and running with local AI on your computer because it is baked into Windows 11 via Microsoft's Copilot+. Like any modern AI using a mixture-of-experts approach, it's a brilliant tool that summarises useful information from a wide variety of resources, saving you the time and effort to do so. Furthermore, Copilot+ can draft emails, blog posts, social media updates, and marketing copy. It can also help brainstorm ideas, refine existing text for tone and clarity, thus making it a digital assistant you never knew you needed but come to appreciate.

The same is true when downloading other local AI programs such as OpenAI's GPT-OSS-20B or Google's gemma3:12. Best done with a comprehensive front-end such as Ollama, getting a local LLM functional, with all its attendant security benefits, is remarkably fast. From first thought of doing so to engaging with local AI takes less than 15 minutes, and remember, running local ensures a seamless, predictable experience without ever having to wait in virtual queues.



SAFE & DEPENDABLE

The biggest credit is that you don't know if the AI is running locally or in the cloud, going by the quality of results for common queries. Of course, local AI is less comprehensive due to its lower parameters - typically 20 billion compared to, say, 500 billion - but there's very little diminution in the perceived quality of the answers. Shows how far and how well-optimised smaller models are in 2025. It's this optimisation that plays into the hands of small devices such as the Sapphire Edge AI PC. What would have been inconceivable just three years ago is reality today.

This case study is particularly pertinent as I am a business owner running an online technology company with members of staff working remotely, so a perfect guinea pig for leveraging the potential benefits of AI in bringing genuinely new ideas to the table. Having used the Sapphire Edge AI mini PC as a daily driver, and swapping between mid-sized LLMs regularly, I'm now a firm believer and convert into using AI for many facets of business.

Easy as AI

Cutting through the hyperbole often associated with new technology, it's intuitive to lean on AI for numerous tasks. I can get a summary of my important emails from Copilot+ thanks to the on-board NPU. Previously, I'm forever searching and getting frustrated with my inbox when I can't find a particular email or thread, knowing that it's important. Now, with Recall, I simply describe the email I'm yearning for, and hey presto, it pops up more often than not. They say that a technology comes into its own when it relieves a pain point; AI certainly does. I also find that summaries work well when I don't want to read through a 30-reply thread, so it's very much a digital assistant for my day-to-day work.

Shifting gears into text-generation LLMs, I find the on-device responses work very well as a research and development tool. I can build up a great, targeted proposal for a potential client by asking local AI the right questions, and as much of my work relates to non-disclosure agreements (NDAs) for upcoming products, knowing that everything stays on my device, and therefore not in the cloud, gives me reassurance I'm taking security - my company's and the client's - absolutely seriously.

The security argument of not wanting to use cloud-based AI is more or less rendered moot by the advances of small yet efficient AI models that work well on tiny boxes like the Sapphire Edge AI mini PC series.

Having delved more deeply into business-based AI than ever before, I'd be loathed not to use it in the future. I find it gratifying I can do all of this on a pint-sized machine that's unobtrusive and quiet during the working day.



THE FUTURE IN THE PALM OF YOUR HAND

AI is here to stay

There's been too much investment, amounting to hundreds of billions of dollars in the last five years, for it to be consigned to the dustbin of history. And it's wonderfully useful if orchestrated properly. The prime ability to be an on-demand digital assistant - and an increasingly smart one at that - plays well across practically all business segments. A natural boon to creatives and marketing types, it's also handy in extracting more efficiency from your day in every instance I can think of.

Rapid evolution of AI software has meant that it's now eminently feasible to run capable solutions right on the device sat next to you, thereby bypassing the cloud and the security concerns that go with having your company's valuable data ferried across the web.

IN CLOSING

Cloud-based models are the last word in performance and insight due to the massive computational power available in huge data centres, but most folk, whether for work or play, actually don't need them. Rather, I'd prefer to run more optimised AI software on my device, where I have more control over what it's doing and how it's connected, and that's readily available right now.

If you've been following the evolution of AI, you'll know it will only get better over time. That's why it is more important than ever to equip your business with easily deployed AI-capable PCs. In that regard, the Sapphire Edge AI mini PC series is a great starting point.