

A Canalys Special Report

Now and next for AI-capable PCs

What impact will the introduction of AI-capable PCs
have on the overall market in 2024 and beyond?

January 2024

Contents

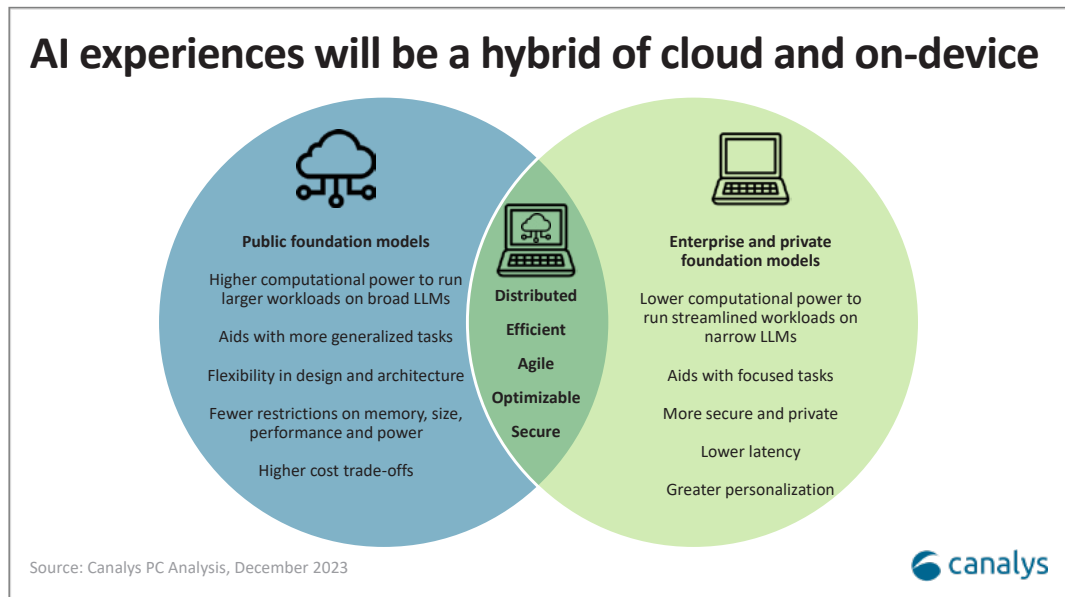
1. Introduction.....	3
2. Impacts of on-device AI on the PC industry	5
What is the definition of an AI-capable PC and the total addressable market?.....	5
Processor vendors lead the evolution of AI-capable PCs.....	7
Key components will upgrade alongside processors.....	8
Microsoft’s vision for hybrid AI makes it an influential leader.....	10
Apple has AI at the core of its Macs.....	10
Apple must prepare for renewed challenges arising from competitors’ on-device AI	13
Consumer and commercial adoption of AI-capable PCs	14
Commercial sector will lead in different phases of AI adoption.....	15
Consumers will expect much more than productivity improvements.....	16
3. Opportunities ahead for AI-capable PCs.....	17
The key PC OEMs are fully committed to AI-capable PCs.....	17
AI-capable PCs will boost current and future replacement cycles.....	18
On-device AI can allow PC vendors to differentiate and monetize meaningfully.....	19
4. Challenges and recommendations	20
Channel optimism around the AI opportunity is growing but buy-in is still needed	20
Convincing channel partners is vital for rapid commercial adoption	21
User resistance to AI must be addressed at every stage of taking products to market.....	22
Creating and showcasing impactful experiences will get customers off the fence	23
AI-capable PCs provide the chance for compelling innovation in the industry	24
Next steps	24

1

Introduction

- Canalys proposes an initial definition of an AI-capable PC as a desktop or notebook with a dedicated chipset or block to run on-device AI workloads. As technical capabilities, use cases, AI tools and customer requirements evolve, additional hardware will need to be considered to identify PCs that are optimized for on-device AI workloads.
- Canalys forecasts that 19% of PCs shipped in 2024 will be AI-capable under its initial definition. This share will ramp up quickly with the tailwind of a refresh cycle, leading to 60% of PCs shipped in 2027 being AI-capable. Commercial adoption is expected to pick up more rapidly from 2024 onwards, and by 2027, 60% of all AI-capable PCs will be shipped to businesses.
- The major Windows PC OEMs are fully committed to AI-capable PCs, with Microsoft acting as an influential industry leader as it drives forward its vision for hybrid AI. Elsewhere, Apple has been shipping Macs that qualify as AI-enabled PCs since 2020 and has been building a strong narrative for on-device AI through an empowered developer community and its focus on device ecosystems and privacy.
- The emergence of AI-capable PCs brings with it several key opportunities for industry players, including a boost to current and future device refresh cycles and the chance to meaningfully differentiate with out-of-the-box AI tools and attached services.
- Key challenges to overcome include getting channel partners on board to drive rapid commercial adoption, overcoming user resistance through transparent and ethical practices, and creating and showcasing meaningful on-device AI experiences to excite the potential user base.

The launch of ChatGPT in November 2022 was a watershed moment, and after rapid growth, generative AI is now positioned as the single most impactful accelerator across nearly every facet of the technology industry. The implications for the smart device space are incredibly important, given that smart devices are the mechanisms through which most users build and interface with AI tools. The role of the PC as the primary device for productivity, especially in business environments, presents a massive opportunity for embedded AI capability to deliver compelling innovation that fundamentally alters the user experience on PCs for the better. As the PC industry braces for a wave of device refreshes four years on from the COVID pandemic and ahead of the end of life of Windows 10, it is vital for players across OEMs, software and operating system providers, component suppliers and the channel to communicate the benefits of AI-capable PCs to customers effectively. Successfully doing so will help reignite a market that has been in an extended period of decline for several quarters, driving higher overall value through the sale of more expensive devices, accelerated hardware refresh cycles and the opportunity to attach AI features and services to customers.



The new era of AI computing will be hybrid. While cloud-based AI tools have driven the initial wave of interest and innovation, the benefits of moving certain workloads onto client devices are quickly becoming apparent. Privacy and security become critical as businesses and consumers increasingly look for AI use cases that leverage models trained on enterprise and personal data. This also provides the following benefits:

- On-device AI functionality helps provide peace of mind that certain private data and insights never leave the PC.
- Local inferencing also means that connectivity issues don't restrict the usability of certain AI tools and that workloads that require low-latency performance can be run anywhere and anytime.
- With the use of AI models set to increase exponentially, associated costs to organizations from accessing cloud resources will ramp up significantly. Moving some workloads to AI-capable PCs will help mitigate this and allow businesses to optimize their use of AI tools according to their budgets.

Fast forward to the future, and a new universal interface might be developed for users to interact with AI models using local data across multiple devices. In that case, users would no longer constrain themselves by using one application sitting on one device to access AI capabilities. AI-capable PCs would become a critical element in the overall sphere of edge AI computing.

2


Impacts of on-device AI on the PC industry

What is the definition of an AI-capable PC and the total addressable market?

With the nascent AI-capable PC category, Canalys initially proposes a narrow hardware-based definition to serve as a starting point for shipment tracking and market-sizing purposes. At a minimum, an AI-capable PC must be *a desktop or notebook possessing a dedicated chipset or block to run on-device AI workloads*. Examples of these dedicated chipsets include AMD’s XDNA, Apple’s Neural Engine, Intel NPU and Qualcomm’s Hexagon Tensor Accelerator.

To align across different industry players’ named offerings, Canalys will refer to this required hardware by the umbrella term of an NPU (Neural Processing Unit).

While almost all PC users are currently able to make use of cloud-based AI applications on their devices, especially those equipped with powerful CPUs and GPUs, the industry needs to ensure that customers can easily identify new products that are specially designed to capture the benefits of on-device AI: real-time responsiveness, low latency, enhanced security and privacy and control over costs. Importantly, they allow users to run AI workloads without an Internet connection. As consumers and businesses increasingly look to supplement the value they get from publicly available, broad-scale generative AI models with finely tuned models based on enterprise and personal data, the benefits of on-device AI will become a key consideration in their device purchasing choice.



Under the current definition, Canalys projects that almost one in five PCs shipped in 2024 will count as AI-capable. These will include all M series Macs, as well as products in the initial launch wave of Windows PCs sporting AI-accelerated chipsets

from the likes of AMD, Intel and Qualcomm. As of now, no significant volumes of Chromebook products are expected to qualify as AI-capable PCs in the short term. But the incorporation of Intel Core Ultra processors alongside appropriate hardware specifications in future releases could see Chromebooks occupy a small niche in the market. The AI-capable share of total PC shipments is expected to rapidly escalate from 2025 as customers are pushed to refresh devices before the end of service for Windows 10 in October of that year.

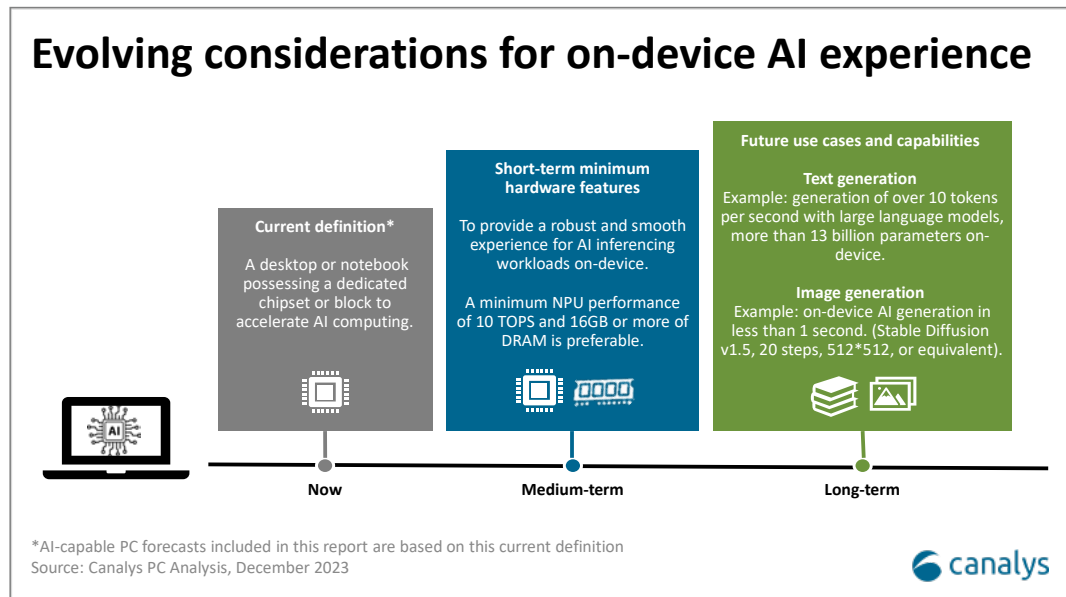
Following initial deployments, both consumers and businesses will have had more time to ascertain how the use of AI tools can supercharge their productivity and deliver compelling user experiences around content creation and personalization. For developers, the breakthrough of on-device AI will encourage the creation of new AI features or applications that harness the NPU for key features or even total use. Canalys expects that demand for bringing these capabilities directly onto their devices will balloon as users seek assurances about security and privacy, smooth use that is not bottlenecked by connectivity and ways of cutting costs. From 2026 onwards, AI-capable PCs are forecast to account for most PC shipments.

Over time, Canalys foresees that its AI-capable PC definition will need to be supplemented with additional considerations:

- As dedicated chipsets or blocks for AI computing become increasingly common in mainstream processors in the coming years, the current definition will become obsolete.
- Additional short-term hardware requirements under consideration include *a minimum NPU rating of 10 trillion operations per second (TOPS) and 16GB or more of dynamic random-access memory (DRAM)*.
- The inclusion of these minimum qualifying specifications will help identify PC products that provide a more robust and smooth experience for on-device inferencing for AI workloads.

As AI applications proliferate over the next few years, this will also set the stage for a more nuanced grading system based on each PC's AI functionalities and capabilities. While most PCs shipped after 2027 are likely to meet Canalys' initial definition, it will be important for customers to be able to distinguish the devices that are best suited to run their specific workloads and that can optimally leverage the suite of AI tools that they choose to invest in. Long-term considerations could include:

- AI-capable PCs could incorporate large language models (LLMs) of a certain parameter size that come pre-installed on the device.
- AI-capable PCs could be benchmarked on the speed of on-device text and image generation. Some benchmarking examples to be considered from 2025 and beyond are listed below, with the specifics likely to evolve over time as on-device AI capabilities improve.



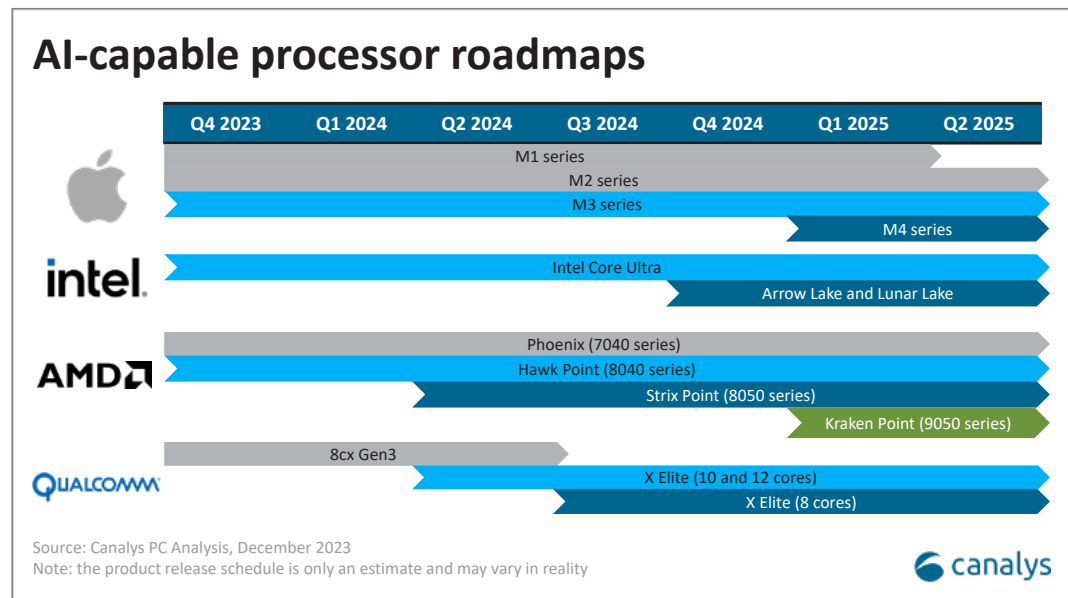
Processor vendors lead the evolution of AI-capable PCs

In the realm of AI-capable PCs, the CPU holds a paramount role. Presently, processor vendors from both the x86 and ARM camps have strategically upgraded their product specifications to align with AI applications and trends in the PC sector. A pivotal shift from previous models is the incorporation of NPUs. While the absence of an NPU doesn't preclude a PC from running AI functions, its inclusion significantly boosts computational efficiency, reduces response time and lowers CPU power consumption. This enhancement positions NPUs as increasingly important components in future PCs.

For companies in the ARM camp, such as Apple and Qualcomm, the integration of NPUs isn't a new concept, largely due to their extensive experience in the smartphone arena. Their in-house technologies currently place them at the forefront in terms of NPU efficiency. Meanwhile, the x86 camp, comprising giants like Intel and AMD, had already identified the significance of this trend years ago. Intel expanded its capabilities through the acquisition of Movidius, while AMD enhanced its portfolio by acquiring Xilinx. These strategic moves are prominently featured in their respective product roadmaps. AMD's recent Phoenix (7040) series, for example, boasts NPUs powered by the XDNA architecture. Intel, conversely, plans to comprehensively integrate NPUs, starting with its Intel Core Ultra series.

In the current landscape of NPU technology, Qualcomm's X Elite stands out with its remarkable 45 TOPS, setting the benchmark as the fastest among all processor manufacturers. Apple's M3 is also a strong contender, achieving a significant 18 TOPS. In the x86 sector, AMD's recent December 2023 release of the Hawk Point (8040) series marks a notable advancement, offering 16 TOPS of computational power. This is a significant leap from its mid-year Phoenix series, which achieved 12 TOPS. Intel, not to be outdone, unveiled its Intel Core Ultra processors in the same month, featuring the Intel NPU and reaching 11 TOPS, which, while impressive, slightly lags AMD's latest product. Despite trailing in TOPS, Intel's deep-rooted influence in the PC market

and superior architectural support will help maintain its competitiveness. Intel has publicly partnered with over one hundred ISVs, which plan to release more than three hundred AI-accelerated features in 2024. Intel’s collaboration with multiple vendors in development planning further bolsters its practical performance against rivals.

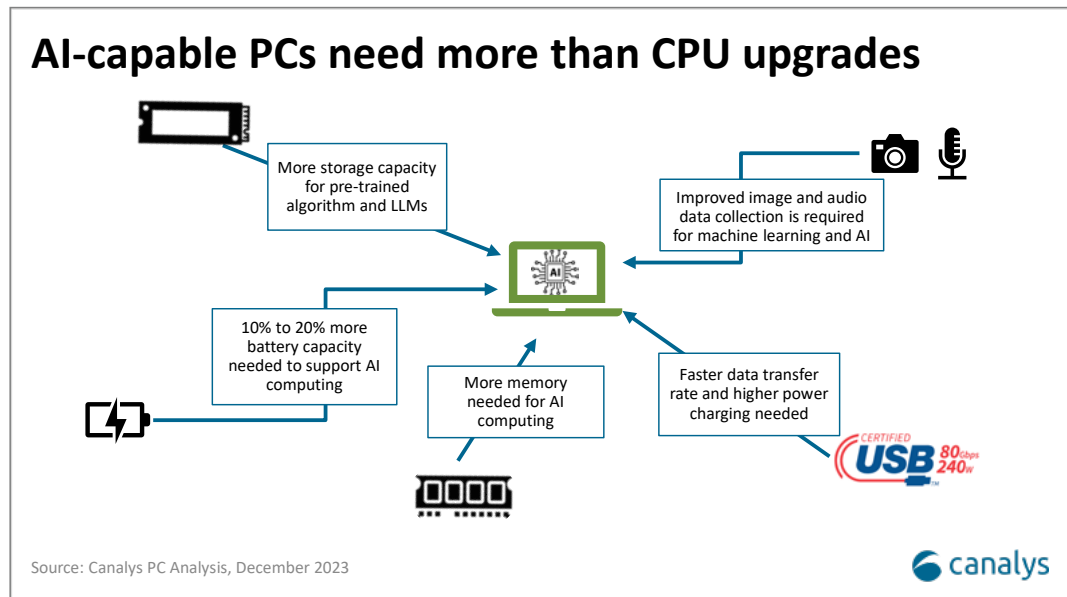


But it’s important to note that the current processor manufacturers do not necessarily use a uniform basis for calculating NPU operations per second. The primary divergence lies in the use of INT4, INT8 and INT16 data structures. Recognizing that smaller data structures can yield higher computational capacities per second is vital, potentially leading to an inflation of performance metrics. Canalys will use INT8 as the benchmark, provided manufacturers publish the relevant data.

As for the availability of these groundbreaking products, MacBook Pros equipped with Apple’s M3 chip and a number of Windows-based products equipped with AMD’s Phoenix processors are already on the market, while Intel Core Ultra devices became available in December 2023. Additional PCs incorporating these advanced processors are slated for release in the first half of 2024. This launch schedule aligns neatly with the expected rollout of a new Windows operating system, marking a significant milestone in the PC industry.

Key components will upgrade alongside processors

Beyond enhancing the CPU, the development of AI-capable PCs also involves major considerations in other hardware areas. The most important element in this context is memory. The intensive data transfer and processing demands of AI applications necessitate memory upgrades in capacity and bandwidth. This trend is expected to significantly elevate the demand for higher-bandwidth solutions with higher density. Concurrently, LPDDR5X and LPDDR5T technology can provide the highest bandwidth among consumer memory products while being energy efficient, which is a strong inducement for PC vendors to increase adoption.



In addition to memory upgrades, many other peripheral component solutions also need to be upgraded to meet the needs of AI-capable PCs. These can be categorized according to their purposes as follows:

- **High-speed data transfer and charging:** AI-capable PCs will need to manipulate data at greater speeds. This is likely to consume more power. We are likely to see a rise in the adoption of advanced USB interfaces, particularly USB4 solutions with higher data speed and power-charging specifications, such as the 80GB and 240W solutions.
- **Equal or improved battery life:** Though processor manufacturers have enhanced power efficiency through the optimization of big and small CPU cores, NPUs and GPU scheduling, the increased computational power demands of AI-capable PCs may lead to reduced battery life, potentially diminishing the consumer experience. Therefore, batteries will need 10% to 20% more capacity than current ones to maintain equal or improved levels of performance.
- **Enhanced real-time sensing and response:** The camera modules and microphones in PCs, which are vital sensing components, lag their counterparts in smartphones in terms of specifications and have shown relatively slow progress. But with growing demand for real-time environmental and user feedback, there is an expectation of enhancements to the number of camera lenses, camera image sensor specifications and the precision of microphones.
- **Storage of multiple trained AI models (LLMs included):** A 50% to 100% increase in storage capacity will be needed to integrate trained AI models, including some LLMs.

Finally, graphic cards (GPUs), widely used in machine learning, high-performance computing and server-side AI computations, have become essential in the AI revolution. In the PC sector, they're often linked to gaming but are sometimes ignored in discussions about AI PCs. But separate GPUs are still vital in AI-capable PCs, especially for intensive

users such as gamers and audiovisual producers, as the integration of CPUs, NPUs and GPUs in processors may not fully meet their demands. While the surge of AI-capable PCs might not drastically change GPU product roadmaps, Nvidia's expected 2025 launch of the GeForce RTX 50 series could be another catalyst for the product segment, especially the gaming division. The collaboration between GPU and processor manufacturers, alongside PC OEMs' expertise in balancing performance with energy efficiency through firmware and software tuning, is essential for unleashing the potential of AI-capable PCs.

Microsoft's vision for hybrid AI makes it an influential leader

Microsoft is perhaps the foremost example of a company betting big on AI to define its future success. On top of its estimated US\$13 billion commitment to OpenAI, it has also invested in companies such as Inflection AI, Adept AI and Builder.ai. While a large part of its AI strategy is focused on providing cloud-based services through Azure, its massive global Windows PC installed base represents a crucial vehicle for delivering new AI experiences to users and convincing developers to create the best AI applications for its ecosystem. As part of this, a significant focus during Microsoft Build in both 2022 and 2023 was the development of hybrid AI experiences spanning cloud-to-device to leverage the benefits of each as use of AI tools proliferates. It has recently made the "Hybrid AI Loop" available, an open-source ONNX Runtime that "supports the same API for running models on the device or in the cloud, enabling hybrid inferencing scenarios where [a developer's] app can use local resources when possible and switch to the cloud when needed."

The following two years represent a crucial window in which successful deployment of AI tools, especially in the workplace, will have profound ramifications for how businesses view their strategic importance. Microsoft's decision in November 2023 to bring Copilot to Windows 10 devices is another strong signal of its commitment to driving user adoption of its native AI tools as part of its AI strategy. Microsoft clearly has a strong desire to migrate its PC user base over to Windows 11 as quickly as possible, with the end-of-service date for Windows 10 looming. But it has also recognized that a slow transition must not be a roadblock to people and businesses incorporating AI experiences into their everyday lives and work. The short-term goal must be to increase penetration and deliver compelling use cases to users with a longer-term view to making the move to future Windows versions exciting by offering new features and significant improvements to AI experiences.

Apple has AI at the core of its Macs

Under Canalys' base definition, Macs have been the only mass-market AI-capable PCs since 2020, when Apple began transitioning from Intel processors to the M1, bringing the Apple Neural Engine (ANE) to the Mac portfolio. The first-generation ANE was released as part of the A11 chip, found in its 2017 flagship smartphone, the iPhone X. Today, every shipping Mac has the Neural Engine.

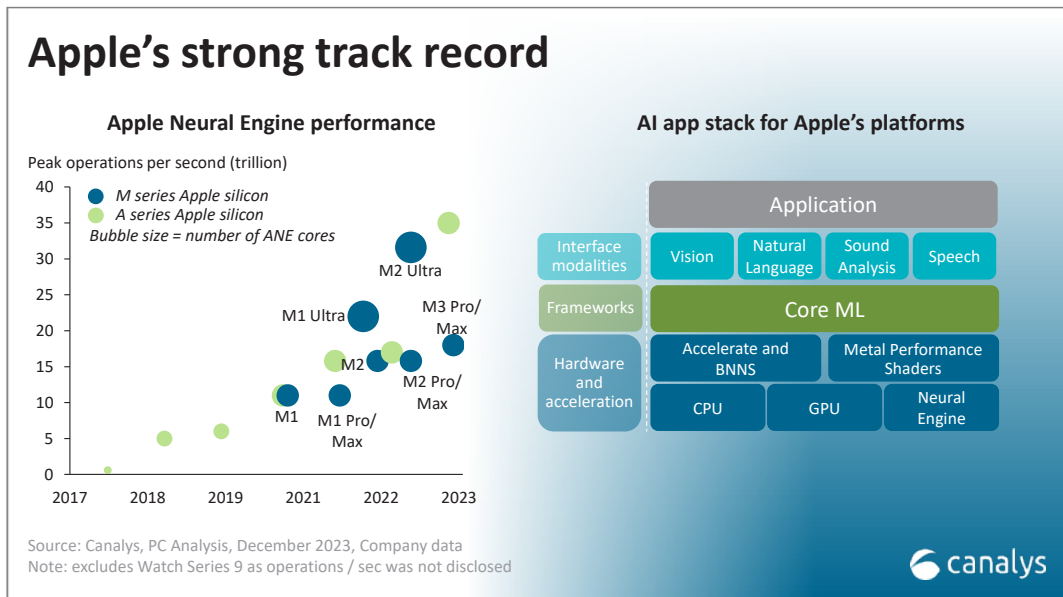
But the company has taken a conservative approach when it comes to incorporating the term "AI" into its product marketing. When asked about the company's apparent lack of

“AI” announcements during its Q4 earnings call, Tim Cook stressed the integral role it has in virtually all Apple products, but that the company tends to “label them as to what their consumer benefit is.” This is true, Apple has integrated many features into macOS that are powered by machine learning.

- The macOS Sonoma update showcases this by enhancing user interaction through machine learning advancements in features such as autocorrect, which now uses transformer language models for more precise text predictions.
- The user experience when collaborating with others is also improved by features such as Center Stage, Portrait mode, Studio Light and adjustable Mic Modes that can filter out background noise. While Translate employs machine learning for accurate language conversion.
- Siri benefits from text-to-speech for natural voice interactions. Voice Isolation distinguishes user speech from background noise during calls, and accessibility features are improved with Personal Voice, allowing users with speech loss to communicate more personally.
- Live Text recognizes text in images and videos and makes them interactable. Visual Look Up uses machine learning to deliver intuitive experiences by allowing users to instantly learn more about the content in a photo, including popular art, world landmarks, plants and flowers, animals, and books.

All these enhancements are made possible with on-device AI, reflecting Apple’s strategy of embedding machine learning deeply into the fabric of macOS to enrich the user experience. When Apple does use AI nomenclature, it’s typically supporting developer outreach. For the last few years, Apple’s Mac marketing strategy around AI has catered to developers. First, Apple provided developers with the tools to create AI-powered apps for their platforms:

- The introduction of the Core ML framework in 2017 made it easy for developers to integrate pre-trained ML models into apps and run them fully on device using CPU, GPU, and ANE. Since then, Apple has continuously expanded its offerings. The Core ML Tools open-source project enables developers to convert a model from popular frameworks like PyTorch and TensorFlow to the Core ML model format and optimize it for performance. More recent additions like on-device ML APIs for Vision, Natural Language, Speech, and Sound Analysis make it easy for developers to integrate powerful Core ML capabilities into their app with just a few lines of code.
- Complementing these efforts, Apple has also continually improved its broader development framework with tools such as Swift and Xcode. The ML Compute Framework aids efficient training on Macs, while advanced APIs allow for the deeper integration of AI. Enhancements to the Metal Performance Shaders and the recent introduction of Dynamic Caching optimize GPU usage, performance, and efficiency. All the while, the company has not wavered in its commitment to user privacy, offering privacy-centric development tools that enable the creation of powerful AI applications without compromising on Apple’s stringent privacy standards.

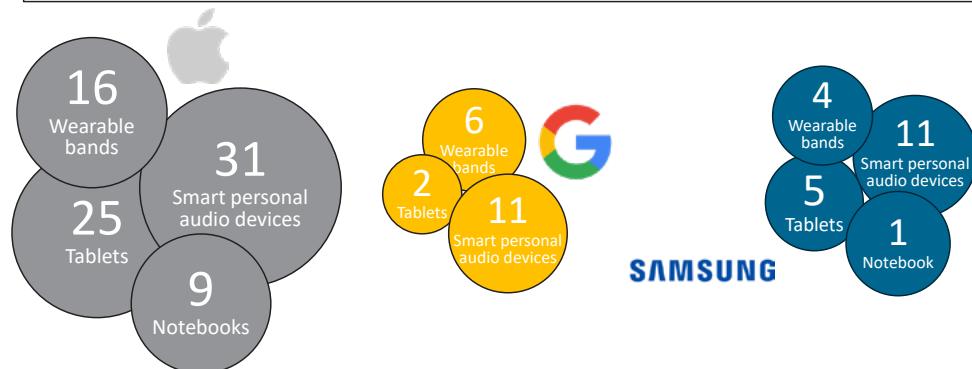


Apple's challenges in the gaming sector provide insight into its strategy for AI. Macs have traditionally lagged in gaming adoption due to a combination of factors. One primary reason is the limited availability of popular gaming titles on macOS compared with Windows, as game developers have often prioritized Windows due to its larger market share and established gaming community. This limitation extends to the gaming experience itself, which can be less optimized on Macs due to differences in graphics technologies and APIs used by game developers, though Apple has been improving on this front recently with chips that feature hardware-accelerated ray tracing, and the launch of the Game Porting Toolkit, and more.

In contrast, when it comes to AI, Apple has actively sought to address such challenges by focusing on developer support and hardware integration. Recognizing that the key to AI success on Mac lies with the developer community, Apple provides comprehensive tools such as Core ML and Create ML. Apple's approach with its unified architecture across its product lines further enhances this, offering a cohesive development environment. This strategy aims to ensure that AI applications on Macs are not only highly functional but also user-friendly and accessible across Apple's device ecosystem.

Apple's ecosystem attach rate remains far ahead

Tablet, IoT and notebook shipments per 100 smartphones shipped worldwide, 2023 year to date



Source: Canalys estimates, Smartphone Analysis, PC Analysis, Consumer IoT Analysis, November 2023
 Note: Apple excludes Beats and Samsung excludes Harman sub-brands, Google excludes Fitbit

The device ecosystem itself is among Apple's largest advantages, and it has an attach rate significantly higher and broader than any other mobility OEM. Its vertical integration approach and signature interoperability between its devices have long been profitable differentiators. When considering on-device AI, Apple's seamless integration across devices becomes even more important and will be a significant advantage when it comes to handling multi-modal AI applications.

Apple must prepare for renewed challenges arising from competitors' on-device AI


Apple has significantly grown its market share among commercial PC deployments in the last several years, catalyzed by the transition to Apple Silicon. Furthermore, a wave of PC refreshes from the Windows version transition is often seen by Apple as an opportunity to make further share gains, especially among businesses. The question facing Apple now is whether Copilot and the slew of Windows-based AI-capable PCs represent a big enough threat to Apple's further growth, especially in the short term. Apple's share gain in the commercial sector did not occur because organizations moved away from Word, Excel and PowerPoint to Pages, Numbers and Keynote. The Mac's value proposition has been highly powerful computers, an increased return on investment and, most importantly, how it fits within a company's existing IT ecosystem. As of today, the Office suite can run natively on Mac and essentially all existing Copilot features are able to run on Mac because the functionality sits in the cloud. But if and when Copilot moves to a hybrid cloud/on-device setup to capture the respective benefits of each, this could change the dynamic. Although this assumes that the trend of collaboration between Microsoft and Apple reverses or slows, and that Copilot would not be brought to a hybrid version on Mac.

Despite its advantages, Apple faces challenges in penetrating a market where Microsoft's Windows is deeply entrenched. Many businesses are heavily invested

in Microsoft’s ecosystem, with established workflows and legacy systems. Apple’s strategy, therefore, might focus not on displacing Microsoft entirely but on carving out niches where its unique strengths in design, user experience and AI capability can be fully leveraged. Apple’s focus is on select verticals, “Pro” users and industries where its products have traditionally been strong, such as creative fields, marketing and software development. These are all fields that will see large benefits from workers upskilling through leveraging AI. These fields have also been identified by Canalys as the early adopters of AI-enabled PCs. Success here will help Apple then expand to other sectors by showcasing the benefits of its AI capabilities and ecosystem integration.

At the core of Apple’s B2B strategy is proposing that organizations give their employees the choice as to how they work, which they are betting will favor Macs in the long run. Catering to end users has worked well, but they aren’t the only stakeholders involved. To combat Microsoft and maintain growth, Apple will need to have a clearer message as to how Macs fit into the landscape of AI-capable PC deployments expected in the coming years. This messaging must be tailored to different decision-makers and partners, from employees to CIOs to the channel and others.

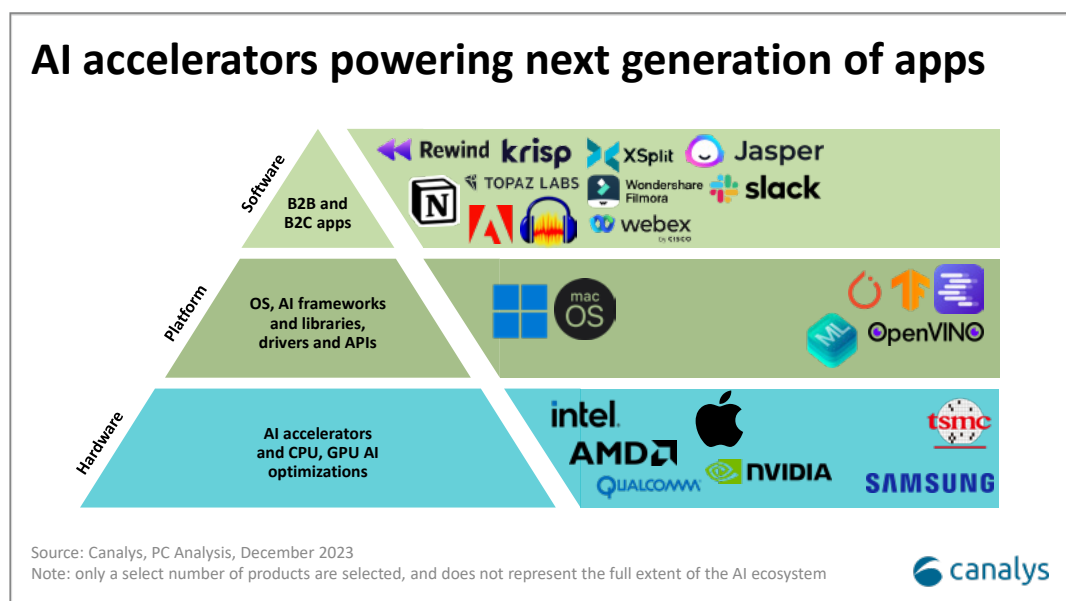
Consumer and commercial adoption of AI-capable PCs



As of Q4 2023, Macs dominate AI-capable PC shipments, as the whole portfolio fits into Canalys’ current definition of the category. But this trend is set to shift as Windows-based products enter the market in the coming year. As Intel, AMD and Qualcomm progressively integrate AI capabilities into their CPU product lines, Windows-based AI-capable PCs from a variety of OEMs will become major growth drivers. The significant improvements in efficiency and productivity are particularly appealing to the commercial sector, especially for tasks such as data processing and content creation. AI-capable PCs are set to become an indispensable part of these operations. Canalys predicts:

- AI-capable PC shipments to the commercial sector will more than triple in 2024 compared with 2023, accounting for 50% of all such devices shipped this year.
- 2025 will bring an inflection point as business demand for AI-capable PCs spikes, leading to more than 100 million devices shipped for the overall category in the year, aided by the peak wave of refresh ahead of the end of service for Windows 10.
- The long-term trajectory for commercial adoption of AI-capable PCs will be positive. In 2027, shipments just to businesses are expected to exceed 100 million devices out of a total of 174 million. More than 60% of all PCs shipped to businesses will be AI-capable, representing a 94% CAGR for the category from 2023 to 2027.
- Though the CAGR for AI-capable PC shipments to consumers will be a more modest 42% over the same timeframe, the category's share of the overall consumer PC market will catch up to the commercial sector in the longer term.

Commercial sector will lead in different phases of AI adoption



Commercial adoption of AI-enabled PCs is set to begin in a very targeted manner. Higher prices and a cost-sensitive market will initially limit these powerful devices to niche roles within organizations. Cost also comes into play when considering the complementary software that companies will use on AI-capable PCs. This selectivity, however, is only the first phase in a broader adoption cycle. As businesses recognize the productivity gains and operational efficiencies enabled by AI, we can expect a shift toward mainstream acceptance. Pioneers in this space, such as research and development professionals, data analysts, and creative designers, will showcase the tangible benefits of integrating AI into their workflows, setting a precedent for broader corporate uptake and IT decision-maker buy-in.

The commercial sector is on the cusp of transformation, with AI-capable PCs at its core. Strategic adoption by tech-savvy businesses will catalyze this change, prioritizing job functions where AI's impact is immediate and pronounced. The future of work will be redefined by AI-driven features that enhance collaboration, streamline content generation and optimize audio and video for a more connected and efficient workplace. As the benefits become apparent and the technology more affordable, on-device AI's role in the commercial space will expand, becoming an integral part of the modern professional toolkit.

The rise of AI-capable PCs will coincide with significant developments in business models. Commercial entities will adopt on-device AI not only for its advanced computational abilities but also for the competitive edge it provides. Integrating AI tools into productivity software and operating systems will pave the way for new business practices and innovation. Furthermore, with the scheduled end of life for older technologies such as Windows 10, businesses will be looking at AI-capable PCs as the next step in their technological evolution, aligning with natural upgrade cycles to modernize their operations.

Consumers will expect much more than productivity improvements

AI-capable PC adoption among consumers is expected to be lower than in the commercial segment in the long run. But consumer interest in generative AI applications is growing and use is becoming more sustained. Users are gradually becoming aware of the array of services on offer and how they can leverage them to improve their computing tasks. For those who have incorporated generative AI into their daily lives, use is robust, with a frequency spanning daily interactions to weekly engagements. To leverage this positive trend in awareness of generative AI and usage behavior leaning toward convenience and ease of use, PC OEMs have an excellent opportunity to provide some differentiated AI features to consumers while operating system providers, such as Microsoft, fully control and unify the user experience around productivity and security features improved by AI.

Performance enhancements, such as efficient task management and power optimization, from the early generations of AI-capable PCs are already expected by users and seen as standard when they buy a new PC (AI-capable or otherwise). PC vendors focusing on the consumer segment must bring out-of-the-box native AI features that benefit consumer use cases, such as creative and professional work, including visual and audio content creation, education and learning. This will be the key to justifying the differences between AI-capable PCs and traditional PCs, which already allow for light-use cloud-based AI productivity workloads. Vendors and other industry players must ensure that consumers have meaningfully transformative experiences in their initial interactions with AI-capable PCs.

Another huge opportunity lies in the device ecosystem, where consumer-focused companies such as Apple have a huge advantage. For consumers, the smartphone will continue to act as the central node in a user's device ecosystem. AI-capable smartphones will act as the hub in the device ecosystem, while AI-capable PCs will be

a vital part of the interconnected experience to leverage data and insights from other edge devices. This will help create more comprehensive solutions for the use cases mentioned above, such as creative work and learning, or even create brand new use cases through interactions with other AI-capable consumer devices.

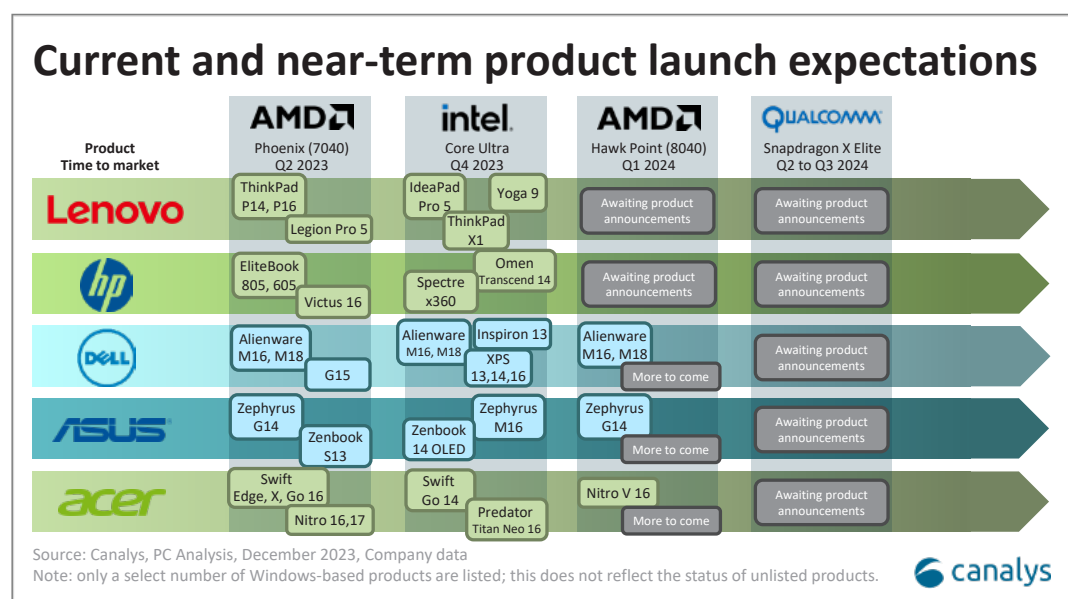
Finally, an important consideration is simply building ease of use and accessibility into the AI experience for consumers. In a future landscape of various AI applications and tools residing on devices across the hardware, operating system and third-party software layers, PC vendors have the advantage of building a simple process for users to access their on-device AI capabilities quickly. A step such as incorporating a physical button that launches the on-device AI assistant or tool that can immediately assist the user and leverage open applications with the correct permissions would help simplify the experience. Additional features such as biometric security or voice recognition could also be included to drive further security and personalization benefits.

3

Opportunities ahead for AI-capable PCs

The key PC OEMs are fully committed to AI-capable PCs

Major PC vendors have publicly signaled that AI-capable PCs will be the most significant upcoming innovation that they are looking to leverage as an accelerator for market recovery in 2024 and beyond. AI PCs were a prevalent theme in the most recent earnings announcement calls of major PC OEMs. Key focuses in the messaging around the category include its creation of “brand new user experiences” (Asus) and its provision of a “huge productivity uplift” (Dell). Vendors are also promoting it as “an inflection point for the industry” (Lenovo) and pointing to the category’s overperformance compared with the total market with “double the expected growth from 2024 to 2026” (HP).



Alongside current Mac products, AMD Phoenix and Intel Core Ultra devices, PC vendors are set to launch a slew of AI-capable PCs in the coming quarters. From a strategic perspective, bringing these new products to market amid the peak wave of Windows refresh will help provide a more compelling reason for customers to transition away from Windows 10. It will also drive users into more premium price bands, with AI-capable PCs expected to be 10% to 15% more expensive than similar models without embedded AI capabilities. But the success of this effort will depend on the vendors working with their key partners (including Microsoft, chipset vendors and B2B and retail channels) to deliver a very clear and unified message about the specific benefits of on-device AI to consumers, businesses, educators and others.

AI-capable PCs will boost current and future replacement cycles

The COVID-19 pandemic brought PCs to the fore for businesses, consumers and schools as the need to bring productivity, education and leisure into the home grew in importance. This led to massive shipment growth across all end-user segments. While remote/hybrid working and learning concepts have outlasted the pandemic, the fallow period for the PC industry over the last two years has shown that additional tailwinds are required to capture the opportunity that has arisen from the increased installed base and positive changes in user behavior. Meaningful innovation has not been easy to come by for the industry over the last few years, leading to a situation where users can hold out on upgrading or replacing their PCs with a minimal impact on the user experience.

Integrating AI capabilities directly into devices provides this innovation, especially with the PC's positioning as the central productivity tool for knowledge workers who will see the quickest gains from incorporating AI tools into their jobs. This will have several positive effects on PC refresh cycles:

- Though most of the momentum from the upcoming refresh in 2024 and 2025 will be driven by the movement away from Windows 10 and from the natural aging of devices, there will be an additional sense of urgency from organizations and consumers who want to be equipped with PCs that can fully capture the potential benefits of AI.
- This will have a modest positive effect on shipment volumes and a more pronounced impact on the average price of PCs bought for upgrades. With successful education and demonstration of the benefits of AI-capable PCs, vendors will be able to move their users onto devices that command higher prices.
- If expectations are successfully met, future replacement cycles will also become more predictable, as users will not be willing to take a step backward in terms of the AI functionality they can access through their PC. Ensuring that user experiences can be significantly improved with next-generation PCs, especially in terms of performance and compatibility with best-in-class AI tools, will strongly encourage organizations to prioritize timely upgrades.

- An example to aim for can be seen in the PC gaming industry, where improvements in hardware to align with the latest game releases help spur quicker refresh cycles than in the general consumer market.

On-device AI can allow PC vendors to differentiate and monetize meaningfully

With Windows OEMs' AI-capable PC roadmaps largely being driven by the availability of next-generation processors from Intel, AMD, Qualcomm and others, there are unlikely to be significant performance differences between devices of similar specifications. While the initial excitement around AI-capable PCs will lead to a case of a rising tide lifting all boats, there will quickly be a need to introduce differentiation to establish a leadership position in the category.

- A key area where PC vendors can set themselves apart will be in the out-of-the-box AI software and applications that they integrate onto devices, including pre-installed LLMs.
- Integrated AI tools with targeted use cases can help position a product for different end-user segments, such as content creators, gamers and "prosumers".

Initiatives such as Lenovo's "AI Innovators Program" and Dell's partnership with open-source LLM platform Hugging Face provide a good foundation. Leveraging these infrastructure-focused steps into the future expansion of on-device AI capabilities will help bring greater customer value and establish vendors as more than just providers of hardware:


- It will help position AI-capable PCs as the anchor around which broader AI services and solutions are delivered, unlocking rich avenues for vendors to monetize the AI opportunity.
- It will be especially important for vendors that are strategically aiming to diversify their businesses away from a high reliance on PC hardware sales.
- Both in-house development of native AI tools and strong engagement with ISVs to deliver rich experiences that are either only available or specifically optimized for their own products will create USPs that have been lacking in the PC industry recently.

4

Challenges and recommendations

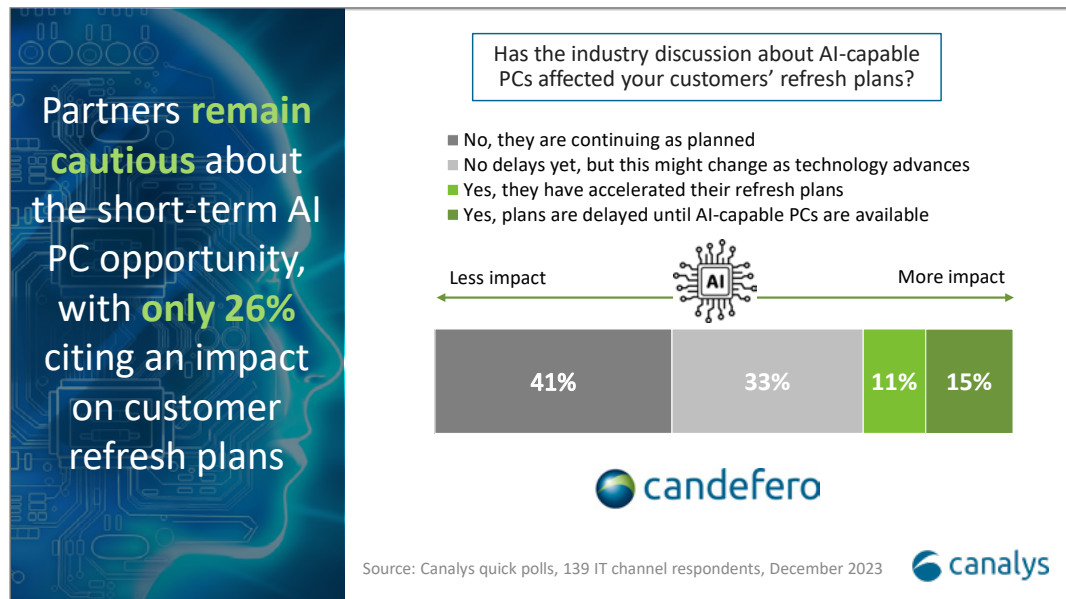
Channel optimism around the AI opportunity is growing but buy-in is still needed

Channel partner optimism for the AI opportunity at large has been on the rise. Polls assessing how much business opportunity AI represents over the next two years show a large spike in positive responses and a drop in neutral responses between 2018 and 2023.



- In all, 41% of channel partners in the most recent survey identify AI as either a “tremendous” or “significant” growth opportunity.
- While a sizable portion of channel partners remained skeptical of the AI opportunity even in 2023, this is to be expected of any emerging technology that does not have immediately clear short-term positive business impacts.
- The important takeaway is that there is momentum in a positive direction, and vendors will have access to a larger pool of partners that are eager to deliver AI-centric products and services to their customers.

But there is currently a contrast in the levels of excitement and hype regarding AI-capable PCs between OEMs and their channel partners. As of the end of 2023, channel partners are still cautious about the short-term opportunity that AI-capable PCs present to drive their customers to refresh devices.



- Of 139 partners surveyed by Canalys in December 2023, only 26% said that announcements around AI-capable PCs had affected their customers' refresh plans.
- 11% stated it has accelerated them and 15% said that they are being delayed until such devices are available.
- Conversations with partners reveal that the underlying value propositions of AI-capable PCs are not sufficiently clear to them or their customers. This partly stems from the fact that, apart from Macs, these devices are not currently available, with 33% of respondents acknowledging that their customers could delay purchase decisions while waiting for technological advances.




Convincing channel partners is vital for rapid commercial adoption


The ecosystem of channel partners (distributors, resellers, systems integrators, managed service providers and others) plays a vital role in the success of any new technology in the commercial sector. Companies such as TD Synnex, Ingram Micro, CDW and downstream partners are at the forefront of delivering hardware, software and services based on their intimate knowledge of their customers' business needs, technical capabilities and budgets. These three considerations will be increasingly important as organizations transition from experimenting with AI tools to developing a focused deployment strategy. This process will quickly increase complexity and cost, especially if AI functionality is brought into work processes that are inappropriate or do not provide a sufficient return on investment. Channel partners with the tools and expertise to guide customers along this new digital transformation journey will act as gatekeepers, so vendors must equip, educate and encourage them to evangelize AI-capable PCs as part of their comprehensive AI solutions.

PC vendors, processor vendors and operating system providers need to treat the channel as a valued partner to drive the growth of the category. Some important near-term considerations are:

- Investment to educate partners on the benefits of AI-capable PCs, both for their own businesses and their customers.
- Empowering partners with market development fund support and detailed knowledge about product roadmaps to allow them to serve customers better.
- Training that involves first-hand experiences of effectively deploying on-device AI in their own organizations.

AI-capable PCs offer host of benefits to partners

<p style="text-align: center; font-weight: bold; margin: 0;">Premiumization</p> <div style="text-align: center; margin: 10px 0;">  </div> <p style="font-size: small; margin: 0;">With advanced hardware specifications and capabilities to deliver productivity gains, AI-capable PC pricing will land 10% to 15% higher than similarly positioned commercial models without on-device AI.</p>	<p style="text-align: center; font-weight: bold; margin: 0;">Faster refresh cycles</p> <div style="text-align: center; margin: 10px 0;">  </div> <p style="font-size: small; margin: 0;">Rapid development in first- and third-party AI tools with on-device features, alongside more frequent hardware improvements, will drive organizations to seek best-in-class experiences for AI-capable PC deployments.</p>	<p style="text-align: center; font-weight: bold; margin: 0;">Attached AI services</p> <div style="text-align: center; margin: 10px 0;">  </div> <p style="font-size: small; margin: 0;">AI-capable PCs will act as an anchor point to provide holistic AI services and solutions, and will play a crucial role in mitigating security and privacy concerns around running AI workloads on sensitive enterprise data.</p>
--	---	---

Source: Canalys PC Analysis, December 2023 

User resistance to AI must be addressed at every stage of taking products to market

To fully leverage the benefits and capabilities of on-device AI for better personalization and efficiency, users will need to give AI models access to important data sitting on their local devices. But this involves treading a delicate line when it comes to user privacy. Companies and individual users are rightfully wary about whether the data collected and stored on their devices is being used by AI models securely and ethically. Another reason for resistance stems from a lack of trust between users and AI solution providers, whether device vendors, operating systems providers or third-party software providers. All AI systems create potentially significant transparency and accountability issues that the tech industry and society must address.

This will be a significant challenge for the entire industry and will take years to overcome. Still, there are clear steps for individual players to take before they bring AI-capable PCs and AI-powered solutions to market:

- Building trust with users by incorporating privacy considerations at every stage of AI product development, from the training of models to how end-user data is accessed. Companies must adhere to high ethical standards throughout this process across all business practices.

- Implementing ethical AI practices to improve transparency, responsibility and accountability. For example, informing users about their data rights and how their information is managed and shared at different customer touchpoints, providing regular audits throughout the AI system's lifecycle, and empowering users to select the amount of data being shared for inference and control the sources during the out-of-the-box AI experience.
- Communicating the above efforts clearly and transparently. Jargon-free communication about privacy policies, data usage explanation, AI performance metrics and any data breaches or misuse. Providing user education on AI technology, data privacy and configuration tools will be essential to improve engagement and limit users putting themselves in harm's way.
- Working with other industry players, governments and regulators to develop frameworks for the use of AI in consumer, business and education settings. As leaders in this area, technology vendors must support these efforts with technical expertise and help promote the democratization of AI.

Creating and showcasing impactful experiences will get customers off the fence

A potential obstacle to adopting AI-capable PCs will be a lack of customer awareness about the breadth of AI experiences available and the specific outcomes they can achieve by incorporating them into their daily lives and work. Organizations will fall on a spectrum of knowledge, openness and technical capability concerning on-device AI, all with the additional consideration of investment budgets. As such, it becomes imperative for PC vendors and their key partners to understand customer requirements, help solve them and then leverage those successes into wider education about how on-device AI can be tailored to other customers' particular needs and constraints. Some ways to demystify AI and build strong cases for its adoption are:

- Engagement with customers across different verticals to understand current digital transformation pain points and suggest appropriate use of on-device AI. This can be achieved both at a detailed level through direct outreach or at larger industry/community-focused events.
- Running pilot programs to produce impactful business or social outcomes, with a focus on the *specific benefits* gained from using AI-capable PCs. Examples could include cost-savings delivered by moving away from cloud workloads, provision of AI functionality in areas with poor connectivity infrastructure or additional data security gained by using on-device AI.
- Published case studies should strongly focus on quantifiable outcomes, for example, in terms of dollars, labor hours or energy savings. They should also drill down into the type of job functions and workloads that can be enhanced by on-device AI so that other potential customers can identify parallels. Specificity and measurability are the keys to reducing confusion about an emerging technology like on-device AI.

AI-capable PCs provide the chance for compelling innovation in the industry

The introduction of AI-capable PCs is a watershed moment for the industry, but effectively capturing the opportunity will require a growth mindset from all players. Firstly, PC vendors must develop business models that stretch beyond hardware sales to deliver differentiated native AI capabilities. Secondly, businesses and users must view AI-capable PCs as part of a larger goal to modernize operations, upskill workforces and participate in the AI revolution through enhanced efficiency and novel experiences. Finally, AI solution providers require an evolving perspective to optimize models for on-device use, enabling ambient computing across edge devices within the broader landscape of hybrid AI experiences.

Fundamentally, AI-capable PCs will only fulfill their potential if industry leaders adopt an open and forward-thinking approach. Vendors must build trust and be transparent around privacy and ethical AI practices to overcome user resistance. Companies need to pilot on-device AI to quantify outcomes and construct business cases to convince partners of the opportunities. Individuals must understand the technology and provide feedback to shape its development responsibly. With collaboration, AI-capable PCs can transform how we learn, work and create for the better, delivering new innovations and capabilities.

Next steps

After exploring the insights in Canalys' report on AI-capable PCs, we invite you to deepen your understanding of this dynamic market.

[Contact us](#) to:

- Arrange a follow-up call with our senior analysts to find out more about the AI PC landscape.
- Learn how our PC Analysis services can benefit your decisions and strategies.
- Get more information on distribution rights for this report.
- Discuss Canalys analysts presenting at your next conference.

The written content of this document represents our interpretation and analysis of information generally available to the public or released by responsible individuals in the subject companies but is not guaranteed as to accuracy or completeness. It does not contain information provided to us in confidence by the industry. Market data contained in this document represents Canalys' best estimates based on the information available to it at the time of publication.

If you wish to use any information from this report in a public forum then you must receive prior explicit written approval from Canalys.

Copyright © Canalys 2024. All rights reserved.
