

Apple's M6 Mac mini and M5 Ultra Mac Studio put local AI at the centre of the desktop Mac

TechRounder PDF Edition

Live article:

<https://www.techrounder.com/apple/apples-m6-mac-mini-and-m5-ultra-mac-studio-put-local-ai-at-the-centre-of-the-desktop-mac/>

By Vipin PG | Published September 16, 2026 | Updated September 26, 2026 | Format: Article | 5 min read

In brief

Apple has refreshed its desktop Mac range with the first 2-nanometer M6 chip and a new M5 Ultra processor with up to 512GB of unified memory. The hardware is aimed at local AI workloads, but the specifications also reveal a wide gap between the affordable Mac mini and the workstation-class Mac Studio.

Quote: Apple has refreshed its desktop Mac range with the first 2-nanometer M6 chip and a new M5 Ultra processor with up to 512GB of unified memory. The hardware is aimed at local AI workloads, but the specifications also reveal a wide gap between the affordable Mac mini and the workstation-class Mac Studio.

Apple has introduced new Mac mini and Mac Studio models with a clear focus on running artificial intelligence on the computer itself. The Mac mini gets the new M6 chip, while the Mac Studio is available with M5 Max and M5 Ultra processors. Both desktops go on sale from September 22.

The announcement is significant for people who want to run local language models, image tools, coding assistants, or automation agents without sending every task to a cloud service. It also changes how the two desktop lines should be compared. The Mac mini is a compact general-purpose computer with stronger AI hardware, while the M5 Ultra Mac Studio is designed around memory capacity and sustained compute.

M6 brings Apple's 2-nanometer process to the Mac

The M6 is Apple's first 2-nanometer chip. It has a 12-core CPU with two super cores, four performance cores, and six efficiency cores. Its 12-core GPU adds a Neural Accelerator to each GPU core, and the chip includes two 16-core Neural Engines.

Apple says the M6 delivers up to 1.2 times the multithreaded performance of the M5 and up to 2.4 times the performance of the M1 in its own selected tests. The company also claims nearly 30 percent higher peak GPU compute for AI than the M5. These figures come from Apple's preproduction hardware testing, so they should be treated as a guide rather than a universal result for every application.

The M6 Mac mini supports up to 32GB of unified memory and up to 170GB/s of memory bandwidth. That is enough for many local AI tools, but memory capacity will limit the size of models it can load. Buyers looking for a desktop that can run larger models locally will need to look beyond the standard M6 configuration.

The Mac mini gets faster, but its role stays practical

Apple claims the M6 Mac mini delivers up to four times faster AI performance, twice as fast storage and graphics, and 40 percent faster CPU performance than the previous model. The company has also added Wi-Fi 7, Bluetooth 6, 2.5Gb Ethernet, and an optional 10Gb Ethernet upgrade.

The new Mac mini starts at \$899 in the United States with 16GB of unified memory. In India, the M6 Mac mini starts at INR99,900, according to CNBC-TV18. Apple India's store lists the new model for availability from September 22.

That combination makes the Mac mini interesting for home servers, developers, small studios, and users who want an always-on machine for local automation. The compact design also keeps the desktop footprint small. The 32GB memory ceiling remains the important buying constraint for anyone choosing it primarily for large AI models.

M5 Ultra is built around memory as much as speed

The M5 Ultra is Apple's highest-end desktop chip. It uses a quad-die design connected through UltraFusion and scales to a 36-core CPU, an 80-core GPU, and 512GB of unified memory. Apple says its memory bandwidth reaches 1.2TB/s, which is 50 percent higher than the M3 Ultra.

The chip also places a Neural Accelerator in every GPU core. Apple claims up to 4.5 times the peak GPU AI compute of the M3 Ultra and up to 40 percent faster graphics performance. It supports hardware-accelerated ray tracing, mesh shading, H.264 and HEVC media engines, four ProRes encode and decode engines, and AV1 decoding.

The memory specification matters more than the core count for some local AI users. A model must fit in memory before the Mac can run it locally, and a larger memory pool can reduce the need for aggressive quantization or offloading. The M5 Ultra does not make every large model fast, but its 512GB configuration gives developers and researchers room to test models that cannot fit on ordinary desktop systems.

Apple is selling the Mac Studio as a local AI workstation

The Mac Studio with M5 Max starts at \$2,499 in the United States, while M5 Ultra configurations start at \$5,499, according to Ars Technica. The price rises sharply with additional memory and storage, so the top specification is aimed at professional users rather than typical desktop buyers.

Apple's description of the Mac Studio focuses on AI researchers, developers, data scientists, visual-effects artists, and other users with workloads that benefit from large local memory. The company says developers can use Core AI, Core ML, Metal, Xcode, and Apple's Foundation Models to run or fine-tune AI models locally.

Ars Technica's assessment is more restrained: the machines are primarily specification updates, and their main change is Apple's decision to market the desktops around local AI inference. That distinction is useful. The new hardware provides more capacity, but the benefit depends on whether the software a buyer uses can take advantage of Apple's GPU, Neural Engine, and unified memory architecture.

What the new lineup means for buyers

The M6 Mac mini is the more sensible choice for coding, media work, everyday desktop use, and smaller local models. Its low starting price and compact size also make it suitable for a home lab or an always-on automation system. Buyers should consider 32GB of memory if local AI is a serious part of the plan.

The M5 Ultra Mac Studio is aimed at a different class of workload. Its main advantage is the ability to keep much larger datasets and models in unified memory while maintaining high memory bandwidth. That can matter for local inference, scientific workloads, complex 3D projects, and large video workflows, but the cost is difficult to justify for ordinary office use.

There is also a practical difference between Apple's performance claims and the experience of running a model locally. Memory capacity determines whether a model fits. Memory bandwidth affects how quickly data moves through the system. Software support determines whether the application can use the available hardware. A buyer who checks only CPU or GPU core counts may choose the wrong Mac.

A strong desktop update with a narrow audience at the top

The M6 Mac mini gives Apple's mainstream desktop a newer chip, faster connectivity, and a stronger local AI story without changing its compact format. The M5 Ultra Mac Studio pushes much further, especially for users who need hundreds of gigabytes of shared memory.

For most buyers, the decision will come down to model size, memory requirements, and software compatibility rather than the headline processor name. The new Mac mini is the practical entry point. The Mac Studio becomes worthwhile when a workload can use its large memory pool and sustained throughput.