

ANSYS DISCOVERY AND AMD EXPAND ACCESS TO HIGH PERFORMANCE SIMULATION FOR MORE ENGINEERS AND DESIGNERS

ANSYS DISCOVERY™ 2026 R1 ON AMD RADEON™ GRAPHICS GIVES ENGINEERING TEAMS A FLEXIBLE, AFFORDABLE PATH TO RUNNING THEIR MOST DEMANDING SIMULATIONS LOCALLY

Thermal and fluid models are not getting smaller, and the pressure to run them early has steadily increased. Many engineering teams now face a persistent mismatch between the highly complex simulations they need to test and the local hardware available to test them. Ansys Discovery 2026 R1 changes that. By extending GPU-accelerated simulation support to AMD Radeon™ graphics cards, Ansys, part of Synopsys, gives engineers and designers more flexibility in how and where they run their most demanding workloads.

Roman Walsh, product manager for Discovery at Ansys, explains, "Our priority at Ansys is to support our customers throughout their workflow, regardless of the hardware or operating systems they choose. We want to meet them where they are."

That means acknowledging the realities of tight IT budgets and hardware scarcity. "As high-performance graphics technology becomes more accessible across the industry, we want to support the systems that IT departments and engineers have available to them," Walsh continues. "AMD brings a very strong offering to this space, and we want to support that growth for our users."

Engineers are no longer locked into a specific hardware purchase because of their simulation software. Their decisions about which platform fits their needs can now be based on a more complete picture of what their entire work environment requires.


BRINGING LARGE-SCALE SIMULATION BACK WITHIN REACH

Using the AMD ROCm™ software stack and HIP C++ programming language, Ansys brought the Discovery 2026 R1 solver to AMD Radeon™ graphics. The solver incorporates libraries optimized for AMD Radeon GPUs, unlocking access to their parallel processing power for large simulation workloads. As a result, Discovery software is no longer tied to a limited set of supported GPU platforms for demanding thermal and fluid workloads.

"For Discovery users, having workstation-level performance for their most demanding problems is a primary driver," Walsh says. "Formerly, high-end simulation meant costly graphics cards or remote machines, and most companies can't justify that expense for all workstations. We frequently see engineers running simulations that require 30-plus gigabytes of GPU memory, but higher-capacity cards at a reasonable price point were not widely accessible. This is particularly apparent in electronics cooling, where the thin, small nature of the components demands exactly the kind of GPU capacity the AMD platform can now deliver."

VARY GRAPHICS MEMORY AS NEEDED TO MATCH HARDWARE CAPACITY TO MODEL COMPLEXITY

AMD Ryzen™ AI Max+ PRO processors employ a unified memory architecture in which the CPU, GPU, and an integrated neural processing unit (NPU) share a large pool of high-speed memory. While the NPU delivers up to 50 TOPS of performance to help accelerate AI tasks, the most relevant benefit for simulation work is variable graphics memory. Users can adjust the amount of dedicated video memory to match the complexity of their models via a simple driver setting. On the HP Z2 Mini G1a workstation, an engineer can allocate up to 96GB of memory directly to the GPU.



The HP Z2 Mini G1a workstation powered by the AMD Ryzen™ AI Max+ PRO 395 processor provides substantial GPU capability in a compact form factor.

Support for variable graphics memory enables 'local bursting,' letting engineers tailor their system's graphics power to the job at hand. Traditional discrete graphics cards have fixed memory, so workloads that exceed that capacity borrow memory across the PCI bus, introducing latency.

"With variable graphics memory, users can scale available GPU memory to match model complexity, helping avoid performance bottlenecks in large simulations," Walsh explains. "We designed Discovery software as an integrated solution that lets designers solve problems naturally. We do not want someone to be limited by their hardware. AMD has a very compelling offering."


AMD Ryzen™ AI Max+ PRO processors help eliminate much of the pre-processing required in traditional simulation, decreasing the need to simplify models to fit within hardware limits. "Simulation users tend to spend a disproportionate amount of their time cleaning geometry, fixing things, and prepping a design," Walsh says. "With the GPU performance these processors deliver, you can skip it. You can have fully defined screws, threads, and components inside your design, and we take care of that automatically. You set up your inputs and hit solve. It is just that simple."

AMD ENGINEERS NOW RUN 30GB SIMULATIONS IN 15 MINUTES

At AMD, thermal engineer Eurydice (Eury) Kanimba relies on those capabilities often. Using Ansys Discovery 2026 R1 with the AMD Ryzen™ AI Max+ PRO 395 processor, she tackles some of the most complex cooling challenges in the semiconductor industry, validating cooling solutions such as cold plates for the next generation of AMD Instinct™ data center GPUs.

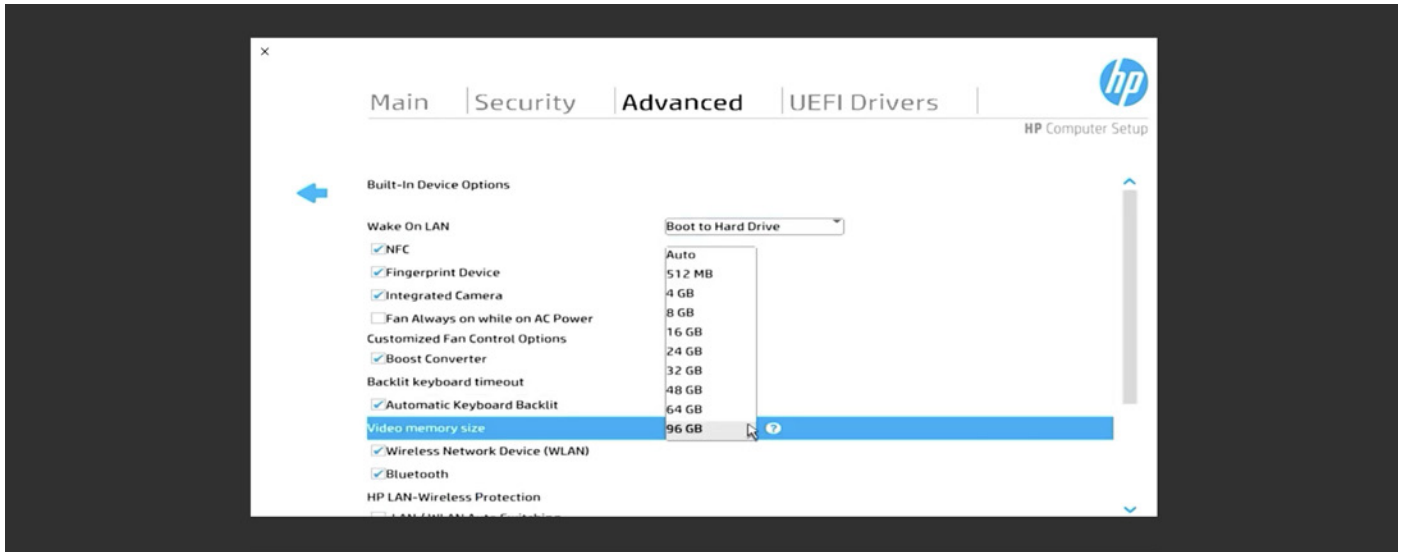
Speed is vital to Kanimba's workflow. She must validate vendor technologies against strict requirements for flow distribution and thermal performance. Because her team evaluates multiple cold plate designs before physical silicon is available, they rely on simulation to understand cold plate performance before manufacturing. This helps them identify possible issues early in the design cycle.

"We need to be ready well before the hardware arrives, so we conduct extensive iterations to qualify vendor designs based on thermal performance and flow distribution," Kanimba says. The result is specific, useful guidance that vendors can act on before a single part is manufactured.



High-fidelity simulation gives engineers deep visibility into internal pressure and flow distributions that are physically difficult to measure in a lab.

The AMD Ryzen™ AI Max+ PRO 395 processor enables Kanimba to perform these evaluations without extensive geometry cleaning. She uses Discovery software's built-in volume-extraction feature to automatically create internal fluid volumes. This works much faster than in other environments, which would have required multiple operations that slowed the process and typically required a more powerful workstation.





Users can easily tune their dedicated video memory pool to match the precise demands of their workflow.

"One major advantage of Ansys Discovery software is the ability to iterate on these designs with ease," Kanimba explains. "Significant curvature often makes meshes difficult to handle in traditional tools, but the features in Discovery replace multiple manual steps and Boolean operations, making the process far more convenient."

The unified memory architecture of the AMD Ryzen™ AI Max+ PRO 395 processor gives Kanimba the flexibility to meet intensive workloads. She can dedicate up to 96GB of memory to even her most demanding models, with room to spare. This capacity allows her to accurately capture the physics of thin structures and sharp edges. Simulations on her HP Z2 Mini G1a workstation now finish in just 10 to 15 minutes as opposed to CPU-based workflows that required hours.

"That speed allows me to explore more options in less time," Kanimba notes. "If I want to adjust a design feature like an orifice opening to meet a target pressure drop, I can do the analysis rapidly and recommend the updated cold plate design to my team and the vendor almost immediately."

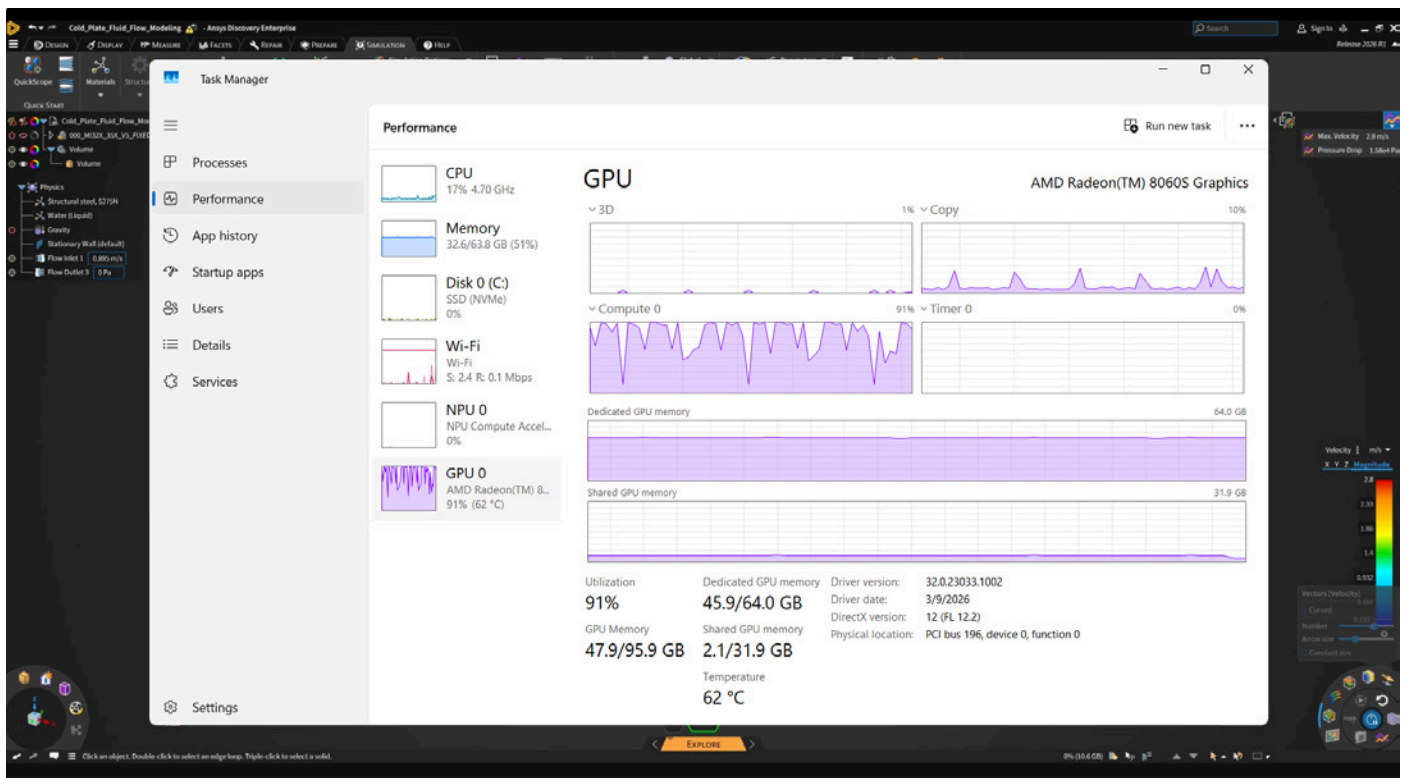


Ansys Discovery™ 2026 R1 allows AMD engineers to import complex CAD models and easily extract internal fluid volumes for rapid velocity simulation.

FINDING DESIGN ERRORS BEFORE THEY BECOME PRODUCTION PROBLEMS

Kanimba's work reveals a clear time-to-decision advantage for local rendering compared with relying solely on cloud tools. By running rapid parametric studies locally, she can vary a feature's radius or adjust a manifold feature to see the immediate flow impact. Her simulations have shown exceptionally high reliability, reaching approximately 90 percent accuracy when compared to test data. That consistency allows AMD to give vendors direction, grounded in simulation data, early in the development cycle.

Kanimba provides a specific example about simulating flow distribution of coolant between the top and bottom plates of a cold plate internal manifold. Because cold plate manufacturing can take 5 to 12 weeks, it's important for both the vendor and AMD to agree on projected thermal and flow performance to ensure timely production and prevent delays. Kanimba explains, "Flow split within a cold plate internal manifold is physically difficult to measure in a lab. Discovery software lets us estimate those conditions and provide the vendor with realistic guidance. The simulation supports our decision to qualify a design for mass production." By catching these issues early in the simulation cycle, AMD helps vendors build the correct version from the start.



The integrated GPU offloads tasks, allowing the workstation to remain responsive during heavy simulation solves.

Beyond simulation speed, the AMD Ryzen™ AI Max+ PRO 395 processor improves daily engineering productivity for her. "The unified memory helps to keep the workstation fully functional during intensive solves," Kanimba notes. "Running heavy simulations on previous systems often caused my computer to freeze, but this platform remains completely responsive. The ability to adjust how much memory I allocate to the GPU and system memory allows me to run a solve in Discovery, participate in meetings, perform data analysis, and use AI tools all at once. That was a challenge on previous workstations. Now it is routine."

AMD AND ANSYS GIVE DISCOVERY COMMUNITY MEANINGFUL HARDWARE CHOICE

By using the AMD ROCm™ software stack and HIP to port its code, Ansys is demonstrating to other developers that the AMD ecosystem is mature and ready for the toughest professional applications.

"The support from the AMD engineering staff has been great," Walsh says. "Getting access to hardware is usually one of our biggest challenges because we need systems available for internal testing. AMD and HP are highly responsive in getting hardware to us, and if our engineers have questions, we receive answers the same day. It is truly a team effort."

The drive to bring Ansys Discovery to AMD graphics platforms has implications that reach well beyond a single discipline. Ansys is actively accelerating solvers across their entire portfolio. The ability to handle larger, more complex problems locally is a significant benefit that spans structural analysis, fluid dynamics, and topology optimization.

For engineering firms, Ansys Discovery software on AMD Ryzen™ AI Max+ PRO processor-based systems change how simulation fits into the design process. Instead of a time-consuming task that interrupts daily work, it becomes something engineers can run, evaluate, and act on in the same sitting. The local memory capacity that variable graphics memory provides makes that possible, even for the largest and most complex models.

With AMD Ryzen™ AI Max+ PRO processors, engineering teams can process large thermal and fluid models directly on their local workstations. This means engineers can quickly deliver crucial simulation data without delay, helping to accelerate the entire design cycle. Extending Discovery software to AMD Radeon™ graphics gives engineers a practical alternative for addressing large models that might otherwise require more expensive GPUs that offer 32GB or more of onboard memory.

"This opens new opportunities for those considering AMD processor-based systems like the HP Z2 Mini G1a," Walsh says. "They can access an enormous amount of GPU capability at an affordable price. The AMD Ryzen™ AI Max+ PRO 395 processor gives them real choices."

With Ansys Discovery 2026 R1, AMD Ryzen™ AI Max+ PRO processors, and HP Z2 Mini G1a and HP ZBook Ultra G1a platforms, engineering teams can bring high-capacity simulation closer to the point of design. That means fewer workflow interruptions, faster design iteration, and more engineers able to use simulation before physical prototypes or vendor lead times slow the process.

DISCLAIMER

The information contained herein is for informational purposes only and is subject to change without notice. While every precaution has been taken in the preparation of this document, it may contain technical inaccuracies, omissions and typographical errors, and AMD is under no obligation to update or otherwise correct this information. Advanced Micro Devices, Inc. makes no representations or warranties with respect to the accuracy or completeness of the contents of this document, and assumes no liability of any kind, including the implied warranties of noninfringement, merchantability or fitness for particular purposes, with respect to the operation or use of AMD hardware, software or other products described herein. No license, including implied or arising by estoppel, to any intellectual property rights is granted by this document. Terms and limitations applicable to the purchase or use of AMD products are as set forth in a signed agreement between the parties or in AMD's Standard Terms and Conditions of Sale. GD-18u.

All performance and/or cost savings claims are provided by the 3rd party organization featured herein and have not been independently verified by AMD. Performance and cost benefits are impacted by a variety of variables. Results herein are specific to such 3rd party organization and may not be typical. GD-181a.

© 2026 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, Radeon, ROCm, Ryzen, and combinations thereof are trademarks of Advanced Micro Devices, Inc. Ansys Discovery is a trademark of Ansys, Inc. Other product names used in this publication are for identification purposes only and may be trademarks of their respective owners. Certain AMD technologies may require third-party enablement or activation. Supported features may vary by operating system. Please confirm with the system manufacturer for specific features. No technology or product can be completely secure.