AMD IT AVOIDS 82% OF DRAM EXPANSION COSTS WITH AMD MEXT™ SOFTWARE

CASE STUDY

AMD IT deployed AI-powered AMD MEXT™ software to add 20% more usable server memory at 18% of the cost of equivalent DRAM capacity



As enterprises scale AI and other memory-intensive workloads, infrastructure decisions must balance performance, cost, deployment requirements, and the flexibility to address what comes next. Server memory has become an increasingly expensive part of that equation. Demand from AI data centers is putting pressure on DRAM supplies, raising costs for businesses that need more memory for their own workloads. In September 2026, TrendForce forecasted another 10–15% quarterly increase in conventional DRAM contract prices for the fourth quarter. For IT teams managing hundreds of servers, even a modest capacity upgrade can now require a substantial investment. AMD IT is addressing this challenge with AMD MEXT™ software, expanding memory capacity without the full cost of physical DRAM.

> **"Rising memory prices are a real challenge for every IT organization."**
> Hasmukh Ranjan, Chief Information Officer, AMD

AMD announced its acquisition of MEXT in June 2026, bringing the latter's predictive memory technology and expertise into the AMD data center portfolio. The company developed an AI-based approach to making flash behave more like DRAM, helping organizations expand usable memory while managing infrastructure costs. AMD is now applying that technology to its own production computing environment.

"Rising memory prices are a real challenge for every IT organization," says Hasmukh Ranjan, Chief Information Officer, AMD. For AMD IT, the challenge was to increase the memory available to engineers without committing the budget that an equivalent expansion using physical DRAM would require. Its production compute infrastructure includes servers based on 4th and 5th Gen AMD EPYC™ Server processors. Adding memory across this fleet represents a significant purchasing decision, with costs multiplying as more systems are upgraded.

OPTIMIZING MEMORY RESOURCES TO MEET GROWING DEMANDS WITH MEXT SOFTWARE

Memory spending also competes with other infrastructure priorities. Every dollar committed to additional DIMMs is a dollar unavailable for internal GPU deployments or other IT hardware investments. AMD IT needed a way to satisfy growing memory requirements while retaining the flexibility to fund those projects.

INDUSTRY
Semiconductor design and manufacturing

CHALLENGES
Expand usable memory across AMD engineering servers while controlling rising DRAM costs and preserving budget for other IT investments.

SOLUTION
Deploy AMD MEXT™ software on AMD EPYC™ processor-based servers, using NVMe flash to extend usable memory capacity at lower cost than equivalent DRAM.

RESULTS
20% more usable memory at 82% lower expansion cost than equivalent DRAM; estimated savings of $2.78 million across ~90 servers.

AMD TECHNOLOGY AT A GLANCE
AMD MEXT™ software
AMD EPYC™ Server CPUs



AMD IT is deploying AMD MEXT™ software across its engineering infrastructure to address growing memory demands and rising DRAM costs.

The question was how to expand usable capacity more economically across infrastructure already supporting AMD engineers. A solution that worked on a few systems also needed to offer a practical route to wider deployment, including both existing production servers and new additions to the fleet. This would be particularly important when many more servers need additional memory at the same time. A small increase per server becomes a much larger purchase across the estate. Reducing the cost of that increment would allow IT teams to consider broader upgrades within the available budget.

> **"MEXT's AI-driven approach gives us a viable way to address cost challenges by expanding usable memory without simply buying more DRAM."**
> Hasmukh Ranjan, Chief Information Officer, AMD

"MEXT's AI-driven approach gives us a viable way to address cost challenges by expanding usable memory without simply buying more DRAM," says Ranjan. AMD MEXT software makes NVMe flash storage appear as an extension of DRAM to the operating system and applications. It moves infrequently accessed memory pages to flash, then uses its AI engine to predict which pages an application will need next. Those pages are brought back into DRAM ahead of the request, helping maintain performance.

The approach gives DRAM and flash different roles within the same memory system. DRAM serves the data needed immediately by the application, while flash provides additional capacity. Predicting demand with the help of advanced AI techniques helps bridge the difference in access speed between the two resources, allowing existing DRAM to support a larger usable memory footprint.

The result is a more optimized memory architecture in which each resource is applied according to the workload's performance, capacity, and cost requirements. Rather than treating every memory need the same, AMD IT can use DRAM where its speed is most valuable and flash where additional capacity can be delivered more economically.

USING AI-POWERED MEXT SOFTWARE TO OPTIMIZE MEMORY ECONOMICS AT SCALE

AMD IT began rolling out MEXT software in September 2026 across production compute infrastructure used by its engineers. This provided an internal test of the technology in the environment it is intended to support, alongside a financial comparison with buying additional DRAM. The deployment reflects AMD's "customer zero" philosophy: using its own IT operations to evaluate emerging technology and develop experience that can inform customers.

The success of this approach must be measured against the cost of adding sufficient physical DRAM to achieve the same memory expansion. AMD IT deployed MEXT software on 86 hosts. Across these systems, the software supplied approximately 58TB of additional capacity. AMD IT calculated that purchasing equivalent capacity as DRAM would have cost approximately $3.3 million. In comparison, the MEXT software deployment cost approximately $601,000.

"AMD IT calculated savings of nearly $2.78 million, or 82%, from its initial MEXT software deployment," says MinQi Bao, Sr. Director, Systems Design Engineering, AMD. In practical terms, the organization could obtain the additional usable capacity for less than a third of the estimated DRAM expenditure.

EXPANDING THE DEPLOYMENT ACROSS AMD

"We have already deployed MEXT significantly in our engineering grid and continue to expand the deployment," says Ranjan. AMD IT plans to install MEXT software by default on new servers while continuing to upgrade existing production systems.



AMD IT's initial MEXT deployment across 86 production servers added approximately 58TB of usable memory capacity, with estimated savings of $2.78 million.

Its internal plan anticipates approximately 700 servers by March 2027, with projected cumulative savings of around $21 million. With a full deployment of 2,000 servers expected by the end of 2027, AMD IT projects savings exceeding $60 million.

Those projections show just how much savings are available from a large fleet. Spending less on memory expansion can free budget for additional IT investments, helping AMD meet engineering infrastructure requirements while supporting further development. The rollout also gives AMD more operational experience to draw on as it evaluates the technology at greater scale.

> **"We have already deployed MEXT significantly in our engineering grid and continue to expand the deployment."**
>
> Hasmukh Ranjan, Chief Information Officer, AMD

The current situation in the wider computing market makes that experience increasingly relevant. TrendForce expects DRAM prices to continue to rise and server supply to remain constrained. Businesses cannot assume that waiting will make their next memory expansion affordable.

For organizations running demanding engineering, analytics, AI, and other memory-intensive workloads, using expensive DRAM more strategically is becoming essential. AMD IT's deployment shows how enterprises can expand capacity on their terms, aligning infrastructure resources with workload needs while preserving budget and flexibility for what comes next.



For organizations running memory-intensive workloads, AMD MEXT™ software offers a more economical way to expand usable memory while preserving budget for other IT investments.

WANT TO LEARN HOW AMD EPYC PROCESSORS MIGHT WORK FOR YOU?
Sign up to receive our data center content:
amd.com/epycsignup

ABOUT AMD

For more than 50 years AMD has driven innovation in high-performance computing, graphics, and visualization technologies. Billions of people, leading Fortune 500 businesses, and cutting-edge scientific research institutions around the world rely on AMD technology daily to improve how they live, work and play.

AMD employees are focused on building leadership high-performance and adaptive products that push the boundaries of what is possible. For more information about how AMD is enabling today and inspiring tomorrow, visit the AMD (NASDAQ: AMD) website, blog, LinkedIn, and X pages.

DISCLAIMERS

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.

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/or 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 set forth in a signed agreement between the parties or in AMD's Standard Terms and Conditions of Sale. GD-18u.

COPYRIGHT NOTICE

© 2026 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, EPYC, MEXT, and combinations thereof are trademarks of Advanced Micro Devices, Inc. Other product names contained herein 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.