MEXT | AMD

WHITEPAPER | Q3'25

Driving Superior Price-Performance for MoonRay with MEXT Predictive Memory™

An Evaluation of MEXT AI-Powered Predictive Memory™ Technology on AMD Ryzen™ Threadripper™ PRO Workstations

---

Executive Summary

This whitepaper showcases MEXT's breakthrough performance on MoonRay, the open-source animation rendering application developed by DreamWorks. Rendering high-fidelity animation frames with MoonRay presents substantial memory and performance challenges, traditionally requiring overprovisioned and costly infrastructure. By integrating MEXT Predictive Memory™ software into the rendering workflow, we achieved near-identical performance to baseline configurations using only HALF the DRAM. These results illustrate MEXT's ability to enable high-performance, cost-efficient rendering at scale.

- 40% reduction in memory-related costs
- 29% reduction in TCO
- 1.4X higher performance-per-dollar
- < 5 min installation, with no changes to underlying hardware

---

Introduction

Modern animation workloads like MoonRay push the boundaries of compute and memory, with high-resolution assets, complex lighting models, and physically accurate rendering techniques driving unprecedented demand for system resources—particularly memory. As studios seek to scale their pipelines, relying solely on traditional DRAM becomes increasingly unsustainable—both technically and economically. The high cost, limited availability, and scaling constraints of DRAM create bottlenecks that threaten production efficiency. Addressing these challenges requires a smarter approach to memory management—one that enables elasticity, optimizes resource utilization, and delivers performance without the penalties associated with overprovisioning.

Hardware Configuration

MEXT-Enabled System

- AMD Ryzen™ Threadripper™ PRO 5995WX, 64-core, 128-Thread 4.5 GHz Processor
- 1TB PCIe Gen 4x4, Gen 5x2 M.2 2280 NVMe Internal SSD
- 512GB Total Memory:
  - 256GB DDR4 RAM
  - 256GB of MEXT Memory™

Control System

- Same workstation as the MEXT-Enabled system
- 512GB Total Memory:
  - 512GB DDR4 RAM
  - No MEXT Memory™

Software Stack

- MoonRay v2.15.0.1
- Rocky Linux 9.5
- Linux Kernel: 5.14.0-503.40.1.el9_5.x86_64
- MEXT Predictive Memory™ software
- MEXT View™ observability software leveraging Open Telemetry (OTel) and Grafana

---

MEXT AI-Powered Predictive Memory™

How It Works

System memory (DRAM) is one of the costliest components involved in modern computing. However, across many business environments, its utilization regularly drops to 50% or below (demonstrated by various studies from leading cloud service providers and hyperscalers).

MEXT's software solution solves this utilization issue by 1) continuously monitoring which memory pages in DRAM actively being utilized, or "hot", and which have gone "cold", 2) offloading the cold memory from DRAM to flash, and 3) leveraging AI to mitigate the effects of flash latency and keep the system performant (via the MEXT Predictive Memory™ Engine).

The MEXT AI-Powered Predictive Memory™ Engine continually predicts which offloaded pages might be requested by the application soon (in other words, which pages are likely to soon go from cold to hot), and transparently moves them back into DRAM before the requests are even made. As a result, the application stays performant because from its perspective, the relevant memory pages are always already resident in DRAM.

Value

This empowers customers to run applications performantly within a much smaller DRAM footprint, with far better utilization of that DRAM footprint—yielding substantially lower costs. This is the key value area that we will focus on in this whitepaper.

In other cases, however, there may be customers whose applications are running out of memory and suffering performance issues as a result. In the past, they would have had to either procure a larger system or shard their data across multiple smaller systems—both of which are complex and expensive paths. With MEXT, customers can cost-effectively expand the effective memory capacity of their system by leveraging flash as memory.

Seamless Implementation

MEXT is a patent-pending, software-only solution that works with any configuration: on-premises or cloud, with any processor, in virtualized / bare-metal / containerized environments, with no changes to the OS or applications. Installation takes less than 5 minutes.

MEXT Solution Components

The MEXT Predictive Memory™ solution consists of 3 primary components: the MEXT Driver, the MEXT Predictive Memory™ Engine, and the MEXT View™ Observability Platform.

MEXT Driver

The MEXT Driver is a dynamically loadable kernel module (which does not alter the standard Linux kernel) that sends process and memory page telemetry data to the MEXT Predictive Memory™ Engine.

MEXT Predictive Memory™ Engine

The MEXT Predictive Memory™ Engine is a user-space process that feeds predictions of which memory pages should be pushed from flash to DRAM—making predictions / inferences in under a fraction of a second. It runs entirely on the local Linux operating system (on a single CPU core) and does not require a GPU.

It was inspired by modern AI techniques based on neural networks. Instead of using these techniques to predict words or natural language patterns (like ChatGPT does), it predicts sequences of future memory page accesses. It consists of a family of models that work together, combining extremely lightweight heuristic predictors with more powerful neural-network models. For any given workload, it automatically adjusts to use the model or group of models that performs best. Continuous observation of which predicted pages were actually used by the application also enables the engine to acquire real-time feedback regarding model accuracy, supporting ongoing adaptation and self-optimization.

MEXT View™ Observability Platform

MEXT also provides a user-space application called MEXT View™ that provides observability / visualization tools to help customers profile their workloads—illustrating how much memory their applications are using at any given time and what portion of this memory is hot / warm / cold. All cold memory pages are good candidates for optimization by MEXT Predictive Memory™ software. MEXT View™ also provides insight into the ongoing prediction accuracy of the MEXT Predictive Memory™ Engine.

Methodology

Observability data was collected from MEXT-enabled workstations during use in a typical work environment over the course of several days. Common memory-intensive activities include creating complex 3D scenes and assets, working with multiple animation layers and camera angles, and previewing and manipulating high quality renders. Resource statistics including CPU, memory and disk I/O utilization were collected. MEXT Predictive Memory™-specific information including prediction accuracy, memory temperature and latency reduction statistics were also captured. Results were assessed qualitatively and quantitatively.

Results

Summary

The MEXT-enabled system performance is almost identical to the control system, and MoonRay application users do not notice a difference in their day-to-day workflows or in the rendering speeds across both setups. The MEXT-enabled system comes with significantly reduced costs (as a result of using only HALF of the DRAM compared to the control). As such, it is clear that MEXT drives outstanding price-performance for MoonRay.

High Performance and Prediction Accuracy

The MEXT Predictive Memory™ Engine keeps the system performant by predicting which pages of memory will be needed soon by the MoonRay application and pushing them from flash to DRAM before they are requested. As seen in the figure below, for the MoonRay rendering job being monitored, the MEXT Predictive Memory™ Engine is able to predict the required memory pages needed by the application with **97.9%** accuracy.

Maximizing DRAM Utilization

As seen in the graph below, the maximum amount of system memory used is 465GB out of 512GB of total system memory; in other words, the workload is using 91% of the total system memory which consists of 256GB of DRAM and 256GB of MEXT Memory™. Without MEXT, this system would likely be non-functional and experience extreme slowdowns or Out-of-Memory (OOM) crashes.

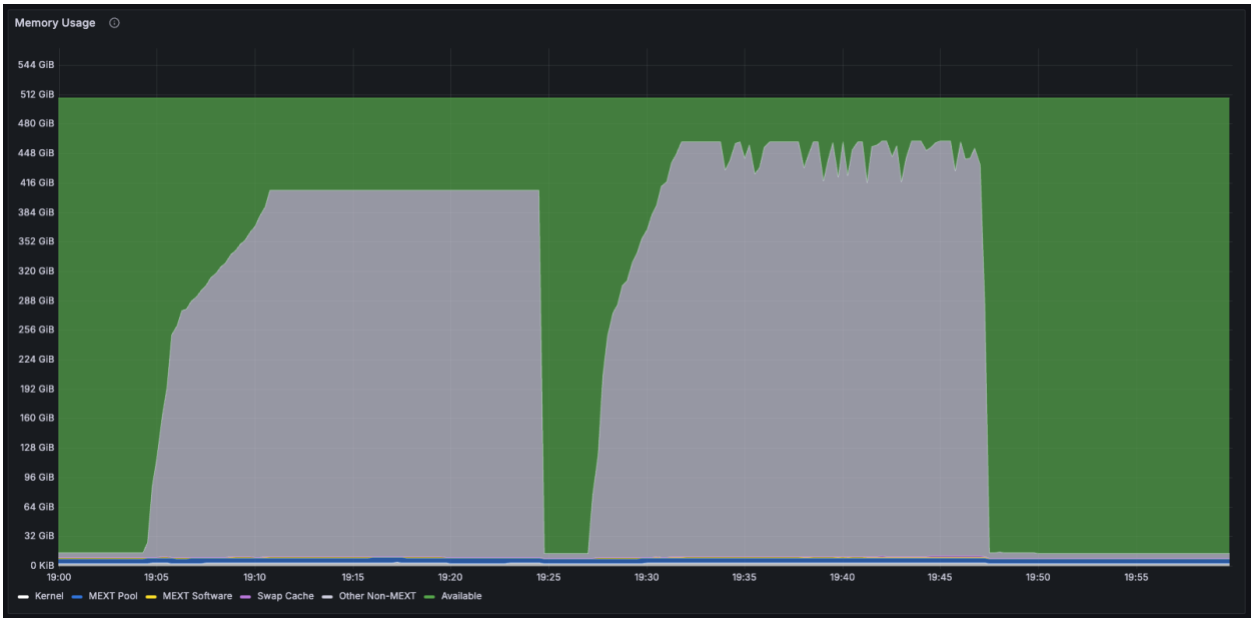

The following Resident Memory Temperature graph examines the actual physical DRAM utilization and color codes it according to how long it has been since the various memory pages in that 15-second interval have been accessed. It is important to note that this graph represents a temporal, not physical over of the system; that is, the memory pages that are hot are not necessarily contiguous and most often are not. This shows that while there are “calmer” periods of memory use, there are also times when 94% of the physical DRAM is being actively utilized, indicating that not only is this system utilizing the DRAM it has, but—as we saw above—it is actively utilizing an almost equal amount of MEXT Memory™.

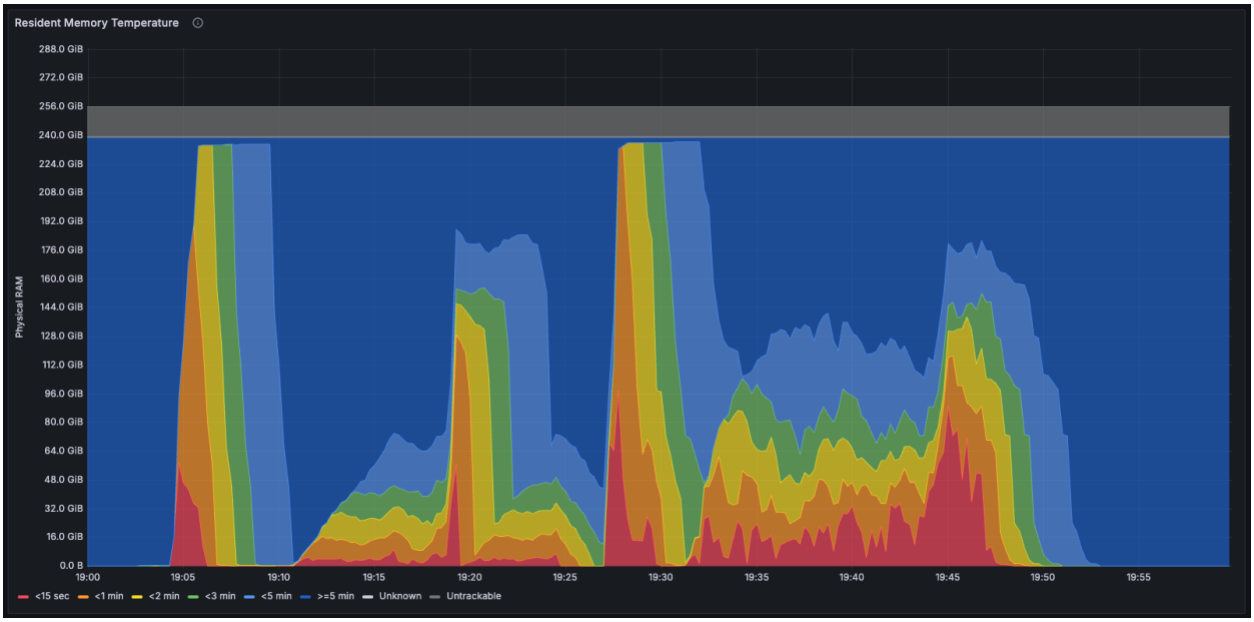

Lower Costs

The ability to reduce the total DRAM by 50% translates to a total cost of ownership reduction of approximately **29%** and a memory-related cost reduction of approximately **40%**. This analysis has been done by comparing the overall hardware cost of the system bills of materials for workstations with 256GB of RAM vs those with 512GB of RAM, while also accounting for the costs of the MEXT Predictive Memory™ solution and the dedicated SSD.

Conclusion

The MEXT Predictive Memory™-enabled system delivers comparable MoonRay performance (maintaining consistent frame quality and render stability) with HALF the DRAM vs the control system. In other words, performance stays roughly the same (let us call it 1X performance) while costs get cut by 29% (0.71X cost), translating to a performance-per-dollar boost of 1.4X.

MEXT software also integrates seamlessly into existing infrastructure and scales effortlessly across studio workstations, making it a practical and powerful upgrade for production environments looking to avoid traditional memory overprovisioning and instead optimize price-performance.

The combination of the leadership processing power of AMD Ryzen™ Threadripper™ PRO processors with MEXT's innovative approach to memory management represents a major transformation in the economics of workstation-based creative production—a highly performant, highly cost-efficient path for studios looking to scale.