

AMD MEXT

AMD MEXT PREDICTIVE MEMORY INSTALLATION GUIDE

October | 2026

Version 1.2

© 2026 Advanced Micro Devices, Inc. All rights reserved.

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's products are as set forth in a signed agreement between the parties or in AMD's Standard Terms and Conditions of Sale.

Trademarks

AMD, the AMD Arrow logo, AMD EPYC, AMD MEXT, and combinations thereof are trademarks of Advanced Micro Devices, Inc. Other product names and links to external sites used in this publication are for identification purposes only and may be trademarks of their respective companies.

* Links to third party sites are provided for convenience and unless explicitly stated, AMD is not responsible for the contents of such linked sites and no endorsement is implied.

Date	Version	Changes
October 8, 2026	1.2	Document Created

Audience

This installation guide describes how to install, configure, and manage AMD MEXT Predictive Memory, in on-premises and cloud environments.

TABLE OF CONTENTS

CHAPTER 1: OVERVIEW 4

CHAPTER 2: System Requirements 5

2.1 - Hardware 5

2.2 - Supported Linux Operating Systems and Kernels 5

2.3 - Security / Access 6

2.4 - Other 6

CHAPTER 3: Deployment Options 7

3.1 - On-Premises 7

3.2 - Cloud 7

CHAPTER 4: Solution Components 8

4.1 - MEXT Driver 8

4.2 - AI Engine 8

4.3 - MEXT Daemon 8

CHAPTER 5: Software Installation 9

CHAPTER 6: Post-Installation Configuration 10

CHAPTER 7: Service Management 11

7.1 - Service Names 11

7.2 - Commands 11

7.3 - Examples 11

CHAPTER 1: OVERVIEW

Optimize price-performance for memory-intensive applications with AMD MEXT Predictive Memory.

Server memory (DRAM) is one of the biggest cost components in the data center, but it's common to find that over half of the provisioned memory isn't actively utilized by applications and can be considered cold. MEXT solves this utilization issue by offloading cold memory to a **50x lower-cost location**—flash—and leveraging AI to mitigate the effects of flash latency.

The AMD MEXT AI Engine continually predicts which offloaded pages might be requested by the application 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 already resident in DRAM. This empowers customers to run applications within a much smaller DRAM footprint—yielding substantially lower costs.

On the other end of the spectrum, there may be customers whose applications are running out of memory and suffering performance issues as a result. In the past, these customers would have 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 and flexibly create a large memory system in a matter of minutes.

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.**

CHAPTER 2: SYSTEM REQUIREMENTS

2.1 - HARDWARE

- **CPU:** 8+ physical cores or vCPUs
  - AMD EPYC™ 7003 Series processors (or newer)
  - AMD Ryzen™ Threadripper 5000 Series processors (or newer)
  - Must provide `avx` (Advanced Vector Extensions) and `constant_tsc` (Invariant Time Stamp Counter) CPU flags; for virtual machines, this may require the virtual processors be set to type `host`.
- **DRAM:** 64GB+
- **120MB** of disk space in the `/opt` filesystem
- **Unused/unmounted PCIe® NVMe™ mixed-use SSD**, any form factor. The drive must be presented to the operating system as a native NVMe device (e.g., `/dev/nvme1n1`).

	Minimum	Recommended
PCIe version	Gen3	Gen4 or higher
Drive Size	1.75x desired MEXT Memory	2x+ desired MEXT Memory
Grade	Consumer	Enterprise
Endurance Rating	1DWPD for Workstations	3DWPD for Servers

Table 1: Solid-State Drive (SSD) Requirements

2.2 - SUPPORTED LINUX OPERATING SYSTEMS AND KERNELS

AMD MEXT Predictive Memory works closely with the Linux kernel to manage application memory placement. This matrix lists the OS distribution and kernel combinations that MEXT regularly tests and fully supports; if your OS/kernel combination is not listed, contact MEXT – many similar combinations run without modification to the software and may be made available for evaluation and future support.

Distribution	OS Version	Minimum Kernel Version
Red Hat Enterprise Linux[1]	RHEL 10.1+	6.12.0-55.9.1
	RHEL 9.7-9.8	5.14.0-611.5.1
	RHEL 9.3-9.6	5.14.0-362.8.1
	RHEL 8.10	4.18.0-553
	RHEL 8.6	4.18.0-372.9.1
Oracle Linux[2]	OL 10.1+	kernel-6.12.0-55.9.1
	OL 9.7-9.8	kernel-5.14.0-611.5.1
	OL 9.3 - 9.6	kernel-5.14.0-362.8.1
	OL 8.10	kernel-4.18.0-553
	OL 8.6	kernel-4.18.0-372.9.1
SUSE Linux (SLES)	SLES 15 SP6-SP7	6.4.0-150600.21.3
	SLES 15 SP4	5.14.21-150400.24.103.1
Ubuntu Linux	26.04	7.0.0-14-generic
	24.04	6.8.0-31-generic
	22.04	5.15.0-43-generic

Table 2: Operating System and Kernel Compatibility Matrix

[1] Rocky Linux support follows the RHEL Support Matrix for Rocky Linux release streams. Customers are responsible for obtaining MEXT OS package dependencies for non-current versions of these release streams as maintainers regularly remove dependency packages once newer versions are available.

[2] Oracle Linux is supported in conjunction with the Red Hat Compatible Kernel (RHCK) only; the Unbreakable Enterprise Kernel (UEK) is not currently supported.

AMD recommends using the latest operating system updates.

2.3 - SECURITY / ACCESS

- **Root user access,** either direct or via sudo, is required for installation and configuration but **not** for ongoing usage.
- **Secure Boot must be disabled:** systems with Secure Boot enabled may block necessary kernel functionality.

2.4 - OTHER

No disk swap devices may be configured on the system.

CHAPTER 3: DEPLOYMENT OPTIONS

MEXT is deployed at the Linux operating system level and, therefore, can be easily deployed in both on-premises or cloud environments in **under 5 minutes**.

3.1 - ON-PREMISES

MEXT is deployed via the installation instructions outlined within this documentation for supported Linux operating systems. MEXT can be deployed on Linux in both bare-metal or virtual machine environments. For virtualized environments, MEXT is deployed at the **guest operating system level and not at the hypervisor level.**

3.2 - CLOUD

MEXT can be deployed in cloud environments using either bare-metal or virtualized infrastructure, provided that the hardware and operating system requirements are met. Since MEXT operates entirely at the Linux operating system level, the method of provisioning hardware in the cloud does not affect its deployment.

CHAPTER 4: SOLUTION COMPONENTS

The AMD MEXT Predictive Memory solution consists of three software components which are provided in a single installation package: the **MEXT Driver**, the **AI Engine**, and the **MEXT daemon**.

4.1 - MEXT DRIVER

The MEXT Driver is a dynamically loadable kernel module that does not alter the standard Linux kernel. It provides process and memory page telemetry to the AI Engine and is compiled at installation time to match the host operating system and kernel version.

4.2 - AI ENGINE

The AI Engine is a user-space process that feeds predictions of which memory pages should be pushed from flash to DRAM. 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.

4.3 - MEXT DAEMON

Once the AI Engine has generated its predictions, they are transmitted to the MEXT Daemon. The MEXT Daemon utilizes these predictions to manage the movement of memory pages between DRAM and flash.

CHAPTER 5: SOFTWARE INSTALLATION

Your AMD MEXT representative will provide you with the solution bundle. Note, running the following commands on an existing installation will result in upgrading the software to the new version.

1. Upload the provided installation file to the system.

2. Untar the solution bundle via the tar command. For example:

```
tar -xvf mext-v1.2.0-0.g56cb6540d.tgz
```

3. Enter into the directory created. For example:

```
cd mext-v1.2.0-0.g56cb6540d.tgz
```

4. Read the End User License Agreement.

```
less MEXT_EULA-*.txt
```

5. Read the README.txt file.

```
less README.txt
```

6. Begin the interactive installation.

```
sudo ./install.sh
```

7. See the additional memory now available to your system (+20% of DRAM by default)

```
free -h
```

ⓘ For unattended installation, kindly refer to the usage information provided via:

```
sudo ./install.sh -h
```

CHAPTER 6: POST-INSTALLATION CONFIGURATION

Configuration parameters are documented with usage, default values, and current settings within the following files:

- **AMD MEXT Predictive Memory :** `/etc/opt/mext/mextd.conf`

Upon initial installation, the system is configured to provide MEXT Memory equal to 20% of physical DRAM by default. To allocate a different amount, modify the `mextMemory` parameter via the `mext-config` command and reboot the server.

For example, on a system with 1TB of physical memory, either of the following will configure the system to provide an additional 40% (410GB) of MEXT Memory after reboot (choose only one):

```
sudo mext-config --set-config mextMemory=40%
sudo mext-config --set-config mextMemory=410G
```

The current status of MEXT software components can be determined by executing:

```
/opt/mext/bin/mext_status.sh
```

To validate that the system is functioning correctly, execute the sample workload:

```
/opt/mext/bin/mext_stress.sh
```

CHAPTER 7: SERVICE MANAGEMENT

The MEXT solution is deployed as a Linux system service and can be managed with standard `systemctl` commands:

7.1 - SERVICE NAMES

- **AMD MEXT Predictive Memory daemon:** `mextd`

7.2 - COMMANDS

- **enable:** allow the service to automatically start at boot time
- **disable:** do not allow the service to automatically start at boot time
- **start:** start the service
- **stop:** stop the service
- **status:** show current service state

Attention: Stopping `mextd` is not permitted as it is the equivalent of removing a memory DIMM from a running system. To stop `mextd`, you must disable the service and reboot.

7.3 - EXAMPLES

To enable and start AMD MEXT Predictive Memory daemon:

```
sudo systemctl enable mextd
sudo systemctl start mextd
```