

6 REASONS USING AMD ROCm™ SOFTWARE FOR AI *IS EASIER THAN YOU THINK*

Much of what once held developers back from using AMD ROCm™ software is no longer the case. Here are six reasons why it's easier than ever to build, run, and scale AI workloads on AMD GPUs, CPUs, and NPUs.

1

PYTORCH® NATIVELY SUPPORTS ROCm

ROCm support code is completely upstreamed into the main PyTorch GitHub repository, making it easier to develop AI applications on AMD hardware. With this integration, you don't need to compile a special fork of PyTorch. Instead, ROCm is treated as a first-class backend option.

2

YOUR ENTIRE INFERENCE STACK RUNS NATIVELY ON ROCm

Serving models at scale is where ROCm really stands out. vLLM®, SGLang, llama.cpp, Ollama™, ONNX® Runtime, and ComfyUI® all run natively on AMD GPUs. Whether you're running a single model endpoint or a high-throughput production service, the inference tooling you already know works out of the box.

3

THE AI ECOSYSTEM HAS MOVED AWAY FROM HARDWARE LOCK-IN

AI development has fundamentally shifted. Open-source frameworks, compilers, and libraries are now built to run across hardware backends, not just CUDA®. ROCm software is built around this open approach, helping you take advantage of open-source innovation.

4

MIGRATING MOST AI WORKLOADS TAKES FEW TO NO CODE CHANGES

Moving from CUDA to ROCm takes less engineering effort than you might expect. Most CUDA-based AI workloads can run on AMD GPUs with minimal — or zero — code changes. With modern frameworks, just install ROCm builds and go. Even with custom dependencies, migrating often means small, localized updates rather than major refactoring.

5

ROCm IS A MATURE, PRODUCTION-READY PLATFORM

ROCm includes the optimized libraries, profiling tools, and multi-GPU support that production AI workloads demand. It ships on a regular release cadence, validated against real enterprise workloads and enabled by active contributions from the open-source community.

6

YOU CAN COLLABORATE WITH AMD EXPERTS AND PEERS

Whether you're evaluating ROCm software for the first time or optimizing production workloads, help is available. Beyond documentation, tutorials, technical guidance, and open-source projects, get community support from AMD experts and technical peers through the free [AMD Developer Program](#) and the AMD Developer Community [Discord](#) and [Discourse](#).

ROCm IS BUILT FOR TODAY'S *AI DEVELOPERS*

The AI ecosystem has evolved. For developers looking to run AI workloads on AMD GPUs, ROCm software offers a mature platform backed by the frameworks and open-source community that power modern AI.

Try out ROCm for yourself with sample AI workloads.

Cloud and Data Center AI

Local AI