PROWESS

The Developer's Toolkit of Choice for Simplifying Local Inferencing

Which toolkit best accelerates your path from a custom AI model to production across any PC hardware: the Intel® Distribution of OpenVINO™ toolkit, the Qualcomm® AI Engine Direct SDK, the Lemonade Server SDK, or the Core ML® framework? Prowess Consulting research, commissioned by Intel, finds out!

The Only Toolkit to Top All Rankings

INTEL® DISTRIBUTION OF OpenVINO™ TOOLKIT

⭐ 5

The Intel® Distribution of OpenVINO™ toolkit received top ratings for hardware support, platform compatibility, model conversion, inference, community, and more across Windows®, Linux®, and macOS®, with growing support for Arm® CPUs.

QUALCOMM® AI ENGINE DIRECT SDK

⭐ 2

LEMONADE SERVER SDK

⭐ 3

CORE ML® FRAMEWORK

⭐ 3

Test Developer's Journey
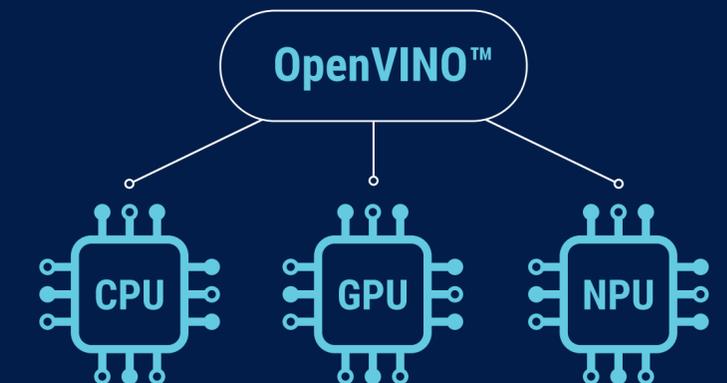
IMPORT → CONVERT/COMPRESS → VALIDATE → TARGET

EVALUATE — PACKAGE — INTEGRATE

Seamless End-to-End Quantization and Deployment

Of the toolkits tested, the Intel® Distribution of OpenVINO™ toolkit was the only one that completed quantization and successfully deployed a chatbot across all models.

OpenVINO™

CPU   GPU   NPU

The Intel® Distribution of OpenVINO™ toolkit made our large language model (LLM) journey smooth and straightforward. Compact, container-ready, and consistent across hardware and operating systems, it delivers a repeatable, developer-friendly pipeline without workarounds. That, combined with its broad hardware and operating system reach, makes the Intel® Distribution of OpenVINO™ toolkit the clear, developer-friendly winner in our testing.

Read the technical research reports:
The Intel® Distribution of OpenVINO™ toolkit vs. the Lemonade Server SDK
The Intel® Distribution of OpenVINO™ toolkit vs. the Qualcomm® AI Engine Direct SDK
The Intel® Distribution of OpenVINO™ toolkit vs. Core ML®