


# OpenVINO™

2023.1

90%

year over year developer download rate<sup>1</sup>

## More Developer Choice



- Automatically import, convert PyTorch and TensorFlow models.
- Use your API of choice based on workload needs.
- Under-the-hood performance with Torch.Compile for StableDiffusion 1.5 and 2.1.

## Gen AI Support



- Expanded model coverage for chatbots, code generation and more.
- Improved GPU support and memory consumption for dynamic shapes.
- More advanced model compression techniques for LLM model optimization.

## Run AI anywhere



- Full support for Intel Meteor Lake, NPU.
- Streamlined AI model deployment via OpenVINO Runtime and Model Server.
- Access to MediaPipe for multipurpose AI.

The Intel Distribution of OpenVINO toolkit makes it easier to **optimize and deploy AI inference at the edge**, enabling developers to write once and bring AI everywhere.

As Intel's **AI inferencing and deployment runtime of choice for developers on Client and Edge platforms**, OpenVINO automatically splits how models are inferred using targeted compute resources. This capability helps to **make generative AI more accessible** for real world scenarios.

**“Ensuring that developers have access to leading-edge tools and capabilities that allow them to write once and deploy anywhere is a key focus for Arm. We are pleased to see initiatives such as OpenVINO that specifically target the deployment of complex AI models to Arm CPUs and enable critical cross-platform support as more intelligent compute moves to the edge.”**

**Dermot O'Driscoll**

*Vice president of Product Solutions, Infrastructure Line of Business, Arm*

<sup>1</sup>Internal Intel Data