



Open Source Software Ready to Maximize Value of New Intel Hardware

Intel contributions to open source projects optimize and demonstrate continued upstream value for frameworks, such as deep learning, machine learning and data prep.

Intel is focused on fostering an open ecosystem to help make programming easier for developers and catalyze community innovation for accelerated computing. The company has a rich history spanning decades of investment in open source software development. Its worldwide network of engineers delivers significant technical contributions to hundreds of independent open source projects, along with providing maintainers and sponsorships, and collaborating with an ever-expanding universe of open source communities.

From the Linux kernel and virtualization software to artificial intelligence (AI) frameworks and widely adopted developer tools, a sampling of open source community projects to which Intel contributes includes AI and machine learning like TensorFlow and PyTorch; cloud native and microservices like Kubernetes, Istio and Envoy; databases including MySQL, RocksDB, PostgreSQL and PrestoDB; data services like Kafka; big data with Spark and Hadoop; networking projects like DPDK and eBPF; programming technologies like LLVM, GCC, OpenJDK, Python and Go; and many others.

Leading up to the launch of the 4th Gen Intel® Xeon® Scalable processors, Intel® Xeon® CPU Max Series and Intel® Data Center GPU Max Series, strategic optimizations were made to open source projects and tools that fast-track application development and deliver business value.

Below are a few highlights. More information can be found at open.intel.com.

Open Source Projects and Tools

- [Linux kernel](#): The Linux kernel is part of the Linux operating system (OS), the world's largest open source OS.
 - Currently available, contributions to the Linux kernel for the 4th Gen Xeon processors were heavily focused on data movement and accelerators for high-performance, demanding workloads.
 - Intel introduced core infrastructure implementing the Compute Express Link standard, support for Intel® Advanced Matrix Extensions (Intel® AMX), as well as the Intel® Data Accelerator Driver to support Intel® Data Streaming Accelerator and the Intel® In-memory Analytics Accelerator. Intel also introduced support for In-Field Scan, which allows circuit-level tests to run on a CPU core for detecting hardware problems not caught by parity or error correcting code checks.
- [TensorFlow](#) and [PyTorch](#): Both TensorFlow and PyTorch are widely used open source machine learning frameworks in the deep learning (DL) arena.
 - Co-developed and open source Intel® Extension for TensorFlow and Intel® Extension for PyTorch deliver optimizations for an extra performance boost on Intel hardware.



- Updates via Intel AMX, a dedicated matrix multiplication engine built into every core of 4th Gen Xeon processors, can deliver up to 6 times higher gen-to-gen DL training model performance according to MLPerf v2.1 industry benchmark for DL training.¹
- **Kubernetes:** The open source project Kubernetes is a system for automating deployment, scaling and management of containerized applications.
 - Intel device plug-ins for Kubernetes enable support for Intel hardware and accelerators in a Kubernetes cluster, including Intel Data Streaming Accelerator and the Intel Analytics Accelerator in 4th Gen Xeon processors. To ensure Kubernetes cluster support with the latest Intel hardware, Intel provides Kubernetes with extensions to expand support for Intel Data Center GPU Max Series.
 - The Intel GPU aware scheduler and GPU plug-in are available, with incremental functionality for both planned through 2023's first quarter, to align with cloud native best practices and optimize performance.
- **OpenJDK:** OpenJDK is the open source and production-ready version of the Java Platform, Standard Edition.
 - Contributions to the most recent version of OpenJDK improve Java performance on 4th Gen Xeon processors, including optimizations for Vector API, scaling, Crypto/Hash/Checksum and Base64 Encode/Decode, as well as enhancements to Java array copy/clear to use the 512-bit vector width instruction.
 - Developers can also benefit from the cross-modify fence optimization for x86 using 4th Gen Xeon "serialize" instruction.
- **LLVM:** LLVM (low level virtual machine) is a collection of modular and reusable compiler and toolchain technologies that make it possible for developers to construct highly optimized compilers, optimizers and run-time environments.
 - Intel contributions to LLVM include the latest instruction set architecture to support FP16 and Intel AMX and performance tuning of C/C++ code for 4th Gen Xeon processors.
 - LLVM 14 enablement for 4th Gen Xeon processors has been available since March 2022.
 - Intel contributions to Clang, the C/C++ compiler front-end, ensure the latest C++ and C standards are supported, such as the C 23, C++ 20 and C++ 23 Language standards.
 - Intel contributed the LLVM SYCL implementation supporting multiarchitecture accelerator compute.
- **GCC:**
 - Intel contributions to GCC/Binutils/Glibc include the latest instruction set architecture to support Intel® Advanced Vector Extensions 512 (Intel® AVX-512) for FP16 and Intel AMX, and performance tuning of C/C++ code for 4th Gen Xeon processors.
 - GCC 12 enablement for 4th Gen Xeon processors has been available since April 2022.
 - Intel drove the vectorization enabling with GCC -O2 and improved the overall SPECINT performance.

¹ See claim [42] at <https://edc.intel.com/content/www/us/en/products/performance/benchmarks/vision-2022/>

Notices & Disclaimers:

Performance varies by use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex. Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Your costs and results may vary.
Intel technologies may require enabled hardware, software or service activation.

About Intel

Intel (Nasdaq: INTC) is an industry leader, creating world-changing technology that enables global progress and enriches lives. Inspired by Moore's Law, we continuously work to advance the design and manufacturing of semiconductors to help address our customers' greatest challenges. By embedding intelligence in the cloud, network, edge and every kind of computing device, we unleash the potential of data to transform business and society for the better. To learn more about Intel's innovations, go to newsroom.intel.com and intel.com.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

The Intel logo, consisting of the word "intel" in a lowercase, sans-serif font, with a small square above the letter "i".