

PRODUCT BRIEF

Intel® Vision Products
OpenVINO™ Toolkit, Intel® Movidius™ VPUs, Intel® FPGAs



THE OPENVINO™ TOOLKIT

STREAMLINES VISION SOLUTIONS ACROSS A ROBUST INTEL® AI SILICON PORTFOLIO

Fast-track enhanced deep learning inference at the edge

"Our partnership with Intel allows us to bring the power of AI to clinical diagnostic scanning and other healthcare workflows in a cost-effective manner."

*—David Chevalier, principal engineer,
GE Healthcare*

"Dell EMC's surveillance labs validation program and joint collaboration with Intel allows us to deliver proven solutions with the top surveillance camera and video software companies."

*—Ken Mills, general manager,
Surveillance and Security Solutions,
Dell EMC*

Edge-to-cloud vision—from deep learning to artificial intelligence

Data is driving a massive paradigm shift in our understanding of how we use and learn from connected systems in a variety of industries—and much of this data relies heavily on vision systems and their ability to capture immense amounts of information. Regardless of whether it's gathered by connected sensors or smart cameras, or whether it's analyzed at the edge or in the cloud, the data from these vision systems is enabling actionable intelligence for businesses everywhere. Consider machine vision-based predictive maintenance at a manufacturing plant, facial recognition software supporting law enforcement in smart cities or transforming consumer retail experiences, or image-based research on the human genome.

Intel® Vision Products integrate advanced software and hardware to capture complex, dynamic visual content from the edge to the cloud, with exceptional richness and accuracy. By delivering data processing flexibility at the edge—both in cameras and on-premise servers—and scalability in the cloud, these solutions are driving next-generation artificial intelligence and analytics, and enabling powerful deep learning inferencing capabilities across various industries.

And while these industries may be diverse in their business offerings, they have a few big things in common: the need for fast access to data at the edge, powerful performance to process and analyze bandwidth-intensive visual data, and the agility and manageability to quickly turn data into insight. This is where Intel's comprehensive AI portfolio comes into play.

A path to vision innovation: building on top of the OpenVINO™ toolkit

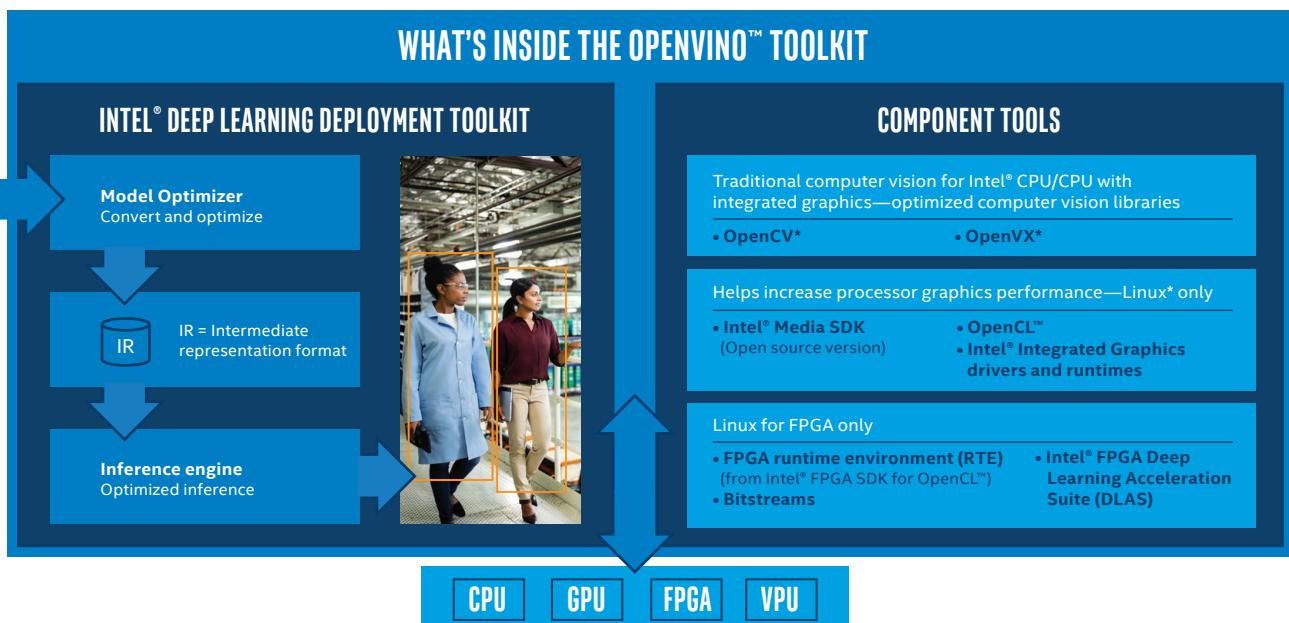
Intel's innovative technology portfolio enables OEMs, ODMs, ISVs, SIs, and developers to speed time to market and deliver increased value to end customers seeking smart vision solutions. The OpenVINO™ toolkit enables acceleration of high-performance computer vision applications and deep learning inferencing. The toolkit helps developers and data scientists speed computer vision workloads, streamline deep learning deployments, and enable easy execution across Intel® platforms, from camera to cloud.

As a primary component of Intel's diverse AI portfolio, the OpenVINO toolkit enhances vision system capabilities and performance through deep learning acceleration silicon—on Intel® architecture (IA) and our Deep Learning (DL) accelerators (FPGA, Movidius™ VPU). This gives innovators added flexibility to choose the right combination of performance, power, and cost efficiency for a specific vision solution.

Benefits of the OpenVINO™ toolkit

- Accelerate application performance across a range of heterogeneous Intel® architectures to deliver fast, efficient deep learning and AI workloads
- Speed up deep learning inferencing from the camera to the cloud
- Offers ease of use and speeds time to market for application developers

You can think of the OpenVINO toolkit as the one suite that enables end-to-end vision solutions—whether you want the flexibility of programmable accelerators or to maximize VPU performance. Applications can be deployed on existing Intel®-based products and infrastructure, while offering the flexibility to integrate efficient, high-performance accelerators without having to use new development methodologies.



Innovation where and when you need it, so you can move forward with AI solutions that deliver edge-to-cloud visual intelligence

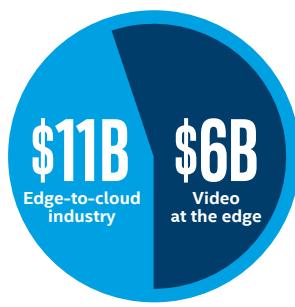
This toolkit enables developers to easily integrate deep learning inference into their applications using industry-standard AI frameworks and standard or custom layers.¹ These can then be deployed across the continuum of Intel-based product lines—from camera to cloud—irrespective of the target platform on which they will be run.² With OpenVINO, developers can write code once and make it future-proof for fast, seamless deployment across current and future Intel® hardware – eliminating application redevelopment.³

The toolkit's Model Optimizer takes care of preparing the algorithm for various Intel® hardware platforms. For example, if the algorithm is originally developed for a camera with an Intel Movidius VPU and later needs to be

redeployed on an Intel® FPGA in a server, the OpenVINO Model Optimizer seamlessly manages code deployment across different target platforms. This allows developers to create the application and choose the Intel architecture compute capabilities necessary for the solution from across the continuum of Intel® technologies—saving significant time and development resources. Using the OpenVINO toolkit, developers write to an API for Intel's inferencing engine.

Intel's innovative solution portfolio is not just about this high-performance software, however; it combines it with industry-leading silicon so that you can more easily make vision-based AI solutions part of your growth strategy.

And the opportunity is compelling.



Up to \$11 billion in edge to cloud industry expenditure is expected by 2021.³ Video captured at the edge by smart devices, sensors, and connected things represents a huge portion—\$6 billion—of this IoT opportunity for everything from smart manufacturing to smart cities.³



Global IP video traffic is expected to be 82 percent of all IP traffic (both business and consumer) by 2021, up from 73 percent in 2016.

This represents a threefold growth and CAGR of 26 percent, with internet video traffic growing fourfold from 2016 to 2021 with a CAGR of 31 percent.⁴

"Hikvision is collaborating with Intel on end-to-end AI/DL solutions, from front-end Movidius cameras to back-end servers. We are excited by the prospect of moving to Myriad™ X. We are also working with Intel on Intel's newly released OpenVINO™ toolkit to achieve higher performance and shorten the development cycle. Hikvision is looking forward to building a strong, long-term relationship with Intel to establish tech leadership in AI/DL-based solutions."

—Dr. Pu Shiliang, chief scientist, Hikvision

TAKE ADVANTAGE OF HARDWARE AND SOFTWARE INTEGRATION FROM EDGE TO CLOUD

OEMS, ODMS, ISVS, SIS

- Scale computer vision and deep learning solutions on existing products and infrastructure
- Integrate efficient, high-performance accelerators

ISVS, DEVELOPERS

- Access vast install base
- Easily port computer vision and deep learning solutions from common software frameworks to Intel® processor and accelerator technologies

VERTICAL INDUSTRY AND ENTERPRISE END USERS

- Get the flexibility to distribute solutions from cloud to edge
- Deliver insights at the right place and time
- Drive results that can transform your business

PUTTING IT ALL TOGETHER

SMART CAMERAS



EDGE COMPUTE



NVRs, gateways, and analytics appliances

DATA CENTER/CLOUD



Integration of high-performance, flexible hardware and software solutions optimized on Intel® architecture helps enable vision innovation

Intel® Movidius™ Myriad™ 2 and Myriad™ X VPUs

These full-fledged system-on-chips (SoCs) are designed specifically for on-device computer vision and neural network inferencing applications. Movidius VPUs have a dedicated architecture offering power efficiency for high-quality image processing, computer vision, and deep neural networks, making them suitable to drive the demanding mix of vision-centric tasks in modern smart devices.

The Intel® Movidius™ Myriad™ 2 VPU offers significant computation performance and image processing capability in a low-power footprint (~1 watt), making it suitable for many mobile IoT end uses. A highly programmable vision processor, the Movidius VPU architecture is built for teraflops of performance supporting multiple cameras with flexible image signal processing pipes per camera and a robust overall dataflow design to mitigate any vision processing bottlenecks.

Target applications include:

- Deep neural network-based classification
- Pose estimation
- 3-D depth
- Visual inertial odometry (navigation)
- Gesture/eye tracking and recognition

The Intel® Movidius™ Myriad™ X VPU is the most advanced Movidius VPU and the first of its kind to feature the Neural Compute Engine—a dedicated hardware accelerator for deep neural network inferences. This accelerator delivers leading performance per watt (one trillion operations per second of deep neural network inferencing⁵), offering a level of on-device performance normally reserved for desktop-class hardware. While the Myriad 2 VPU offers deep neural network support at low power, the Myriad X VPU can reach 10x higher performance for applications requiring multiple neural networks running simultaneously.⁵

The Movidius Myriad X VPU also features 16 configurable MIPI lanes capable of supporting up to eight HD-resolution RGB sensors via direct connections. The high-throughput inline ISP ensures streams are processed at high speeds, and new hardware encoders offer 4K support for both 30 Hz (H.264/H.265) and 60 Hz (M/JPEG) frame rates. And with 2.5 MB of homogeneous on-chip memory, the Movidius Myriad X VPU allows for up to 400 GB/sec of internal bandwidth, minimizing latency and reducing power consumption by minimizing off-chip data transfer.

Intel® FPGAs

Intel® FPGAs offer a wide variety of configurable embedded SRAM, high-speed transceivers, high-speed I/Os, logic blocks, and routing. Built-in intellectual property, combined with outstanding software tools, can lower FPGA development time, power, and cost.

At the edge, these versatile Intel FPGAs provide the flexibility to program and reprogram the rate of acceleration as needs evolve, while providing low latency implementation within a small envelope. Intel FPGAs can combine AI algorithms with other functions into your system in a single chip, for low latency, high throughput, and low total cost of ownership. Because Intel FPGAs are reprogrammable, IoT designers have the flexibility to respond to changing conditions (rapid changes in sensor technologies that require new interfaces, HDR signal processing for image capture in mixed and dark lighting conditions, support for a variety of computer vision algorithms) and implement multiple behaviors within one edge device. Also, due to their parallel computing nature, FPGAs provide accelerated performance for surveillance and other video-enabled applications.

OpenVINO toolkit

The OpenVINO toolkit is a free, downloadable toolkit that helps developers and data scientists fast-track development of high-performance computer vision and deep learning into vision applications. It enables deep learning on hardware accelerators and easy heterogeneous execution across multiple types of Intel® platforms. It includes the Intel® Deep Learning Deployment Toolkit with a model optimizer and inference engine, along with optimized computer vision libraries and functions for OpenCV* and OpenVX*.

This comprehensive toolkit supports the full range of vision solutions, speeding computer vision workloads, streamlining deep learning deployments, and enabling easy, heterogeneous execution across Intel® platforms from edge to cloud. In combination with Intel's diverse AI portfolio, the OpenVINO toolkit provides the power to scale computer vision solutions. The wide range of advanced silicon allows solution providers to match the performance, cost, and power-efficiency required at any node in an AI architecture.

ACHIEVE THE OPTIMAL BALANCE OF POWER EFFICIENCY AND HIGH PERFORMANCE

USE INTEL® MOVIDIUS™ MYRIAD™ 2 VPU FOR ...

- High-performance VPU
- Ultralow-power architecture
- 12 programmable SHAVE cores
- Small footprint for embedded applications

USE INTEL® MOVIDIUS™ MYRIAD™ X VPU FOR ...

- Dedicated neural compute engine
- 16 high-performance SHAVE cores
- Enhanced ISP with 4K support
- New vision accelerators, including stereo depth
- Support for multiple-VPU configuration

Sample use cases

Intel® IoT vision technologies combine end-to-end intelligence and exceptional performance to help improve safety, operational efficiency, productivity, healthcare, and quality of life. Here are just a few of the real-world deployments based on Intel vision solutions.

SMART CITIES

 **Traffic monitoring:** Intel FPGAs and Movidius VPUs on cameras can capture data and automatically send it to traffic intersections downstream for traffic optimization and planning. This information can be directly communicated with drivers via in-car systems or apps for route planning. License plate recognition (LPR) on moving vehicles alerts law enforcement, while identification of available parking spaces aids drivers in congested urban areas.

 **Public safety:** With Myriad VPUs and algorithms developed using the OpenVINO toolkit, trained deep neural networks now have the inferencing capabilities to analyze and identify missing children through facial recognition. Cities equipped with the technology can be alerted when these trained datasets match a face in a crowd with a reported missing child, immediately alerting law enforcement agencies.

INDUSTRIAL MANUFACTURING

 **Industrial automation:** Intel vision solutions are helping enable intelligent factories to converge OT and IT—transforming industrial business models and growth strategies. Solutions based on AI technology produce more detailed, accurate, and meaningful digital models of equipment and processes for nuanced, data-informed planning. Production controls become autonomous and fluid, driving faster time to market.

 **Machine vision:** Augmenting industrial machine vision with AI is enabling more accurate factory-automation applications. Solutions are combining cameras, computers, and algorithms to analyze images and videos for actionable intelligence at the edge. Intel FPGA-powered machine vision supports essential industrial functions, including proactive equipment maintenance, defect detection, quality control, yield management, inventory verification, product labeling, security monitoring, and worker training.

RETAIL

 **Responsive retail:** Retailers using Intel computer vision solutions at the edge can quickly recognize specific customers or customer behavior patterns and deliver personalized, precision marketing. Fast, efficient processing by in-store solutions based on Intel Core processors and Intel FPGAs delivers sales information to associates in near-real time, empowering them with revenue-driving tools and the ability to promote relevant products and services.

 **Operations management:** With Intel® AI technology, in-store vision systems can alleviate the time constraint that many retailers face as data being captured by video feeds makes its way to a distant data center. By using a computer vision solution based on Intel architecture, retailers can streamline operations, manage inventory, optimize the supply chain, and improve merchandising. In turn, they can get more value from their data through scalability, increased efficiency and effectiveness, and lower total cost of ownership (TCO).

Your path to smart vision innovation

Accelerate high-performance computer vision, AI, machine learning, and deep learning solutions from camera to cloud, using Intel's end-to-end solution portfolio.



Learn more about smart video from edge to cloud

Explore Intel IoT smart video solutions at intel.com/smartzvideo.

Download the free OpenVINO™ toolkit at intel.com/visionproducts.

Download the Myriad™ Development Kit (MDK) ›

For more information about Intel® IoT Technology and the Intel® IoT Solutions Alliance, please visit intel.com/iot.

1. The broad set of popular frameworks and topologies supported include, but are not limited to, Caffe*, Caffe2*, MXNet*, Neon*, TensorFlow*, Theano*, and Torch*.
2. While any standard algorithm will run on any Intel® silicon architecture, performance may vary from one architecture to another. In some cases, extra work may be needed using the OpenVINO™ toolkit to port an algorithm from one architecture to a different architecture.

3. "A Strategist's Guide to the Internet of Things," Frank Burkitt, strategy-business.com/article/00294?gko=a9303.

4. cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/vni-hyperconnectivity-wp.html

5. Maximum performance based on peak floating-point computational throughput of Neural Compute Engine. Actual results on deep neural networks may achieve less than peak throughput. Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at intel.com/iot.

Cost reduction scenarios described are intended as examples of how a given Intel-based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

Intel, the Intel logo, Arria, Intel Atom, Intel Core, Intel Movidius, Intel Movidius Myriad, OpenVINO, and Xeon are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.

*Other names and brands may be claimed as the property of others.

© Intel Corporation

OpenCL and the OpenCL logo are trademarks of Apple Inc. used by permission by Khronos.

0518/DW/CMD/PDF 337634-001US