

INCREASE DEEP LEARNING MODEL PERFORMANCE WITH INTEL® MOVIDIUS™ VPU

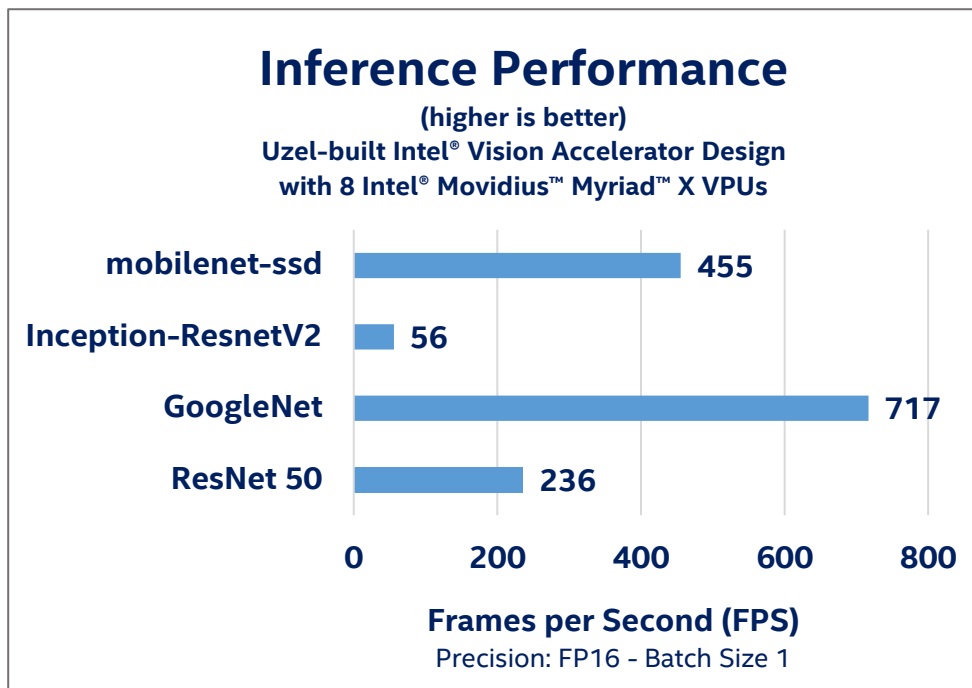
Inference Performance Benchmarks for High-performance Image Throughput

[Intel® Distribution of OpenVINO™ toolkit¹](#) helps accelerate deep learning inference across a variety of Intel® processors and accelerators. Below are performance benchmarks that show public neural networks optimized by the toolkit, that demonstrate performance, value, and flexibility for quick deployment on [Intel® Vision Accelerator Designs with Intel® Movidius™ Vision Processing Units \(VPUs\)](#).

Intel's Movidius VPU is a specialized processor designed to deliver high-performance machine vision and deep learning with compute efficiency. The VPU is ideal for deep learning inference in smart cameras, edge appliances such as network video recorders (NVR), or analytics appliances. These performance benchmarks used an Uzel-built Intel Vision Accelerator Design with 8 Intel® Movidius™ Myriad™ X 2485 VPUs.

THROUGHPUT

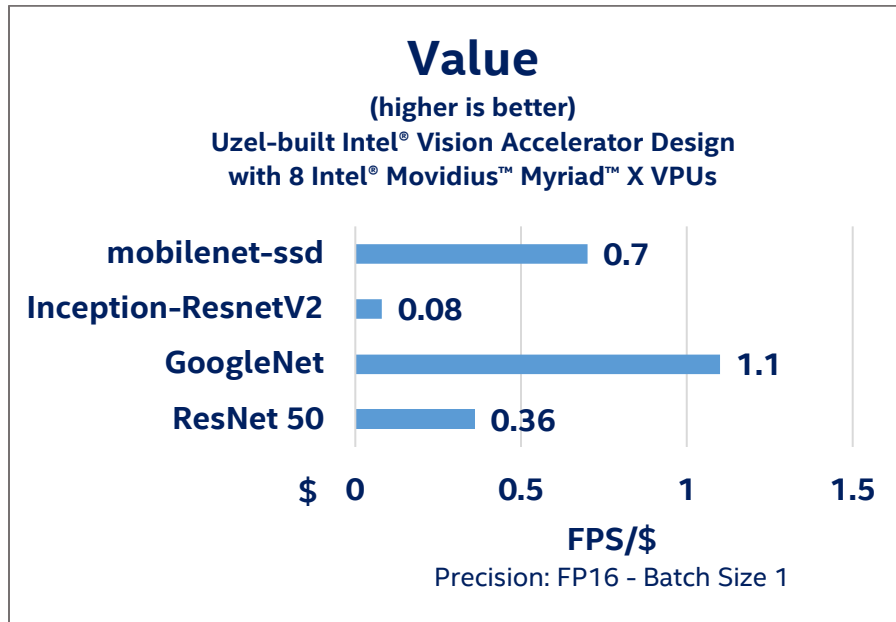
When deploying a system with deep learning inference, select the lowest frame rate that meets solution requirements. Then optimize for power and select analytics algorithms that deliver the best tradeoff between latency and power for the price and performance needed.



Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at intel.com, or from the OEM or retailer. Performance results are based on testing as of March 29, 2019 and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks. CONFIGURATIONS: Testing by Intel as of March 29, 2019. See end of document for configuration details.

VALUE

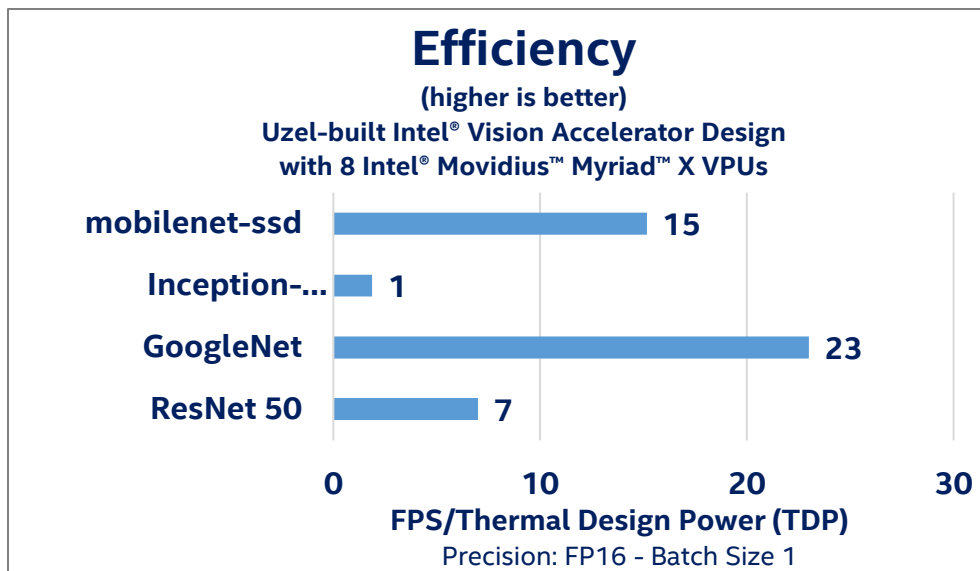
While throughput is important, what is more critical in edge AI deployments is the performance efficiency or performance-per-cost. Application performance in frames-per-second (FPS) per dollar of system cost is the best measure of value.



Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at intel.com, or from the OEM or retailer. Performance results are based on testing as of March 29, 2019 and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks. CONFIGURATIONS: Testing by Intel as of March 29, 2019. See end of document for configuration details.

EFFICIENCY


System power is a key consideration from the edge to the data center. When selecting deep learning solutions, power efficiency (FPS/watt) is a critical factor to consider. Intel designs provide excellent power efficiency to run deep learning workloads.



Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at intel.com, or from the OEM or retailer. Performance results are based on testing as of March 29, 2019 and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks. CONFIGURATIONS: Testing by Intel as of March 29, 2019. See end of document for configuration details.

TOTAL BENEFIT

Combining the factors of value and efficiency can be a good way to compare which hardware yields the best performance per watt and per dollar for your particular use case.



Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at intel.com, or from the OEM or retailer. Performance results are based on testing as of March 29, 2019 and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks. CONFIGURATIONS: Testing by Intel as of March 29, 2019. See end of document for configuration details.

¹An open source version of [OpenVINO™ toolkit](#) is also available.

NOTICES AND DISCLAIMERS

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at intel.com, or from the OEM or retailer. Performance results are based on testing as of March 29, 2019 and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure. Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information, see [Performance Benchmark Test Disclosure](#).

CONFIGURATIONS: Testing by Intel as of March 29, 2019

PRODUCT	UZEL INTEL® VISION ACCELERATOR DESIGN PRODUCTS
PRECISION	FP16
BATCH SIZE	1
PRODUCT TYPE	8 x Myriad X
MOTHERBOARD	Uzel info / US-E1300
LLC CACHE	6144
NUMBEROF SOCKETS	1
HT	OFF
NUMBEROF CPUS	4
SCALING DRIVER	acpi-cpufreq
MODEL	Intel® Core™ i5-6600 CPU @ 3.30GHz

TURBO	ON
PRODUCT	Intel® Core™ i5-6600 CPU @ 3.30GHz
BENCHMARK	12497
THREAD(S) PER CORE	1
SPEEDMIN (MHZ)	3300
SPEEDMAX (MHZ)	3300
STEPPING	3
SPEEDMIN (MHZ)	1000
SPEEDMAX (MHZ)	1000
SLOT	0
SIZE (MB)	16384
VENDOR	029E
SPEED (MHZ)	2133
PARTNUMBER	CMSO16GX4M1A2133C15
TYPE	DDR4
CHANNEL	ChannelA-DIMMO
SLOT	1
SIZE (MB)	0
VENDOR	N/A
SPEED (MHZ)	N/A
PARTNUMBER	N/A
TYPE	N/A
CHANNEL	N/A
SLOT	2
SIZE (MB)	16384
VENDOR	029E
SPEED (MHZ)	2133
PARTNUMBER	CMSO16GX4M1A2133C15
TYPE	DDR4
CHANNEL	ChannelB-DIMMO
SLOT	3
SIZE (MB)	0
VENDOR	N/A
SPEED (MHZ)	N/A
PARTNUMBER	N/A
TYPE	N/A
CHANNEL	N/A
VENDORID	0x8086
VENDOR	Intel® Corporation
NAME	Intel® Gen9 HD Graphics NEO
DRIVER	19.04.12237
VERSION	OpenCL* 2.1 NEO
DEVICE	0
TYPE	GPU
RELEASE	43364
VERSION	5.12
VENDOR	American Megatrends Inc.
KERNEL	4.15.0-29-generic
NAME	Ubuntu* 16.04.6 LTS
LISTED TDP	30 W
LIST OR ODM PRICE	\$650

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice. [Notice revision #20110804.](#)

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

Copyright © 2019, Intel Corporation. All rights reserved. Intel, the Intel logo, Pentium, Xeon, Core, VTune, OpenVINO, are trademarks of Intel Corporation or its subsidiaries in the U.S. and other countries.