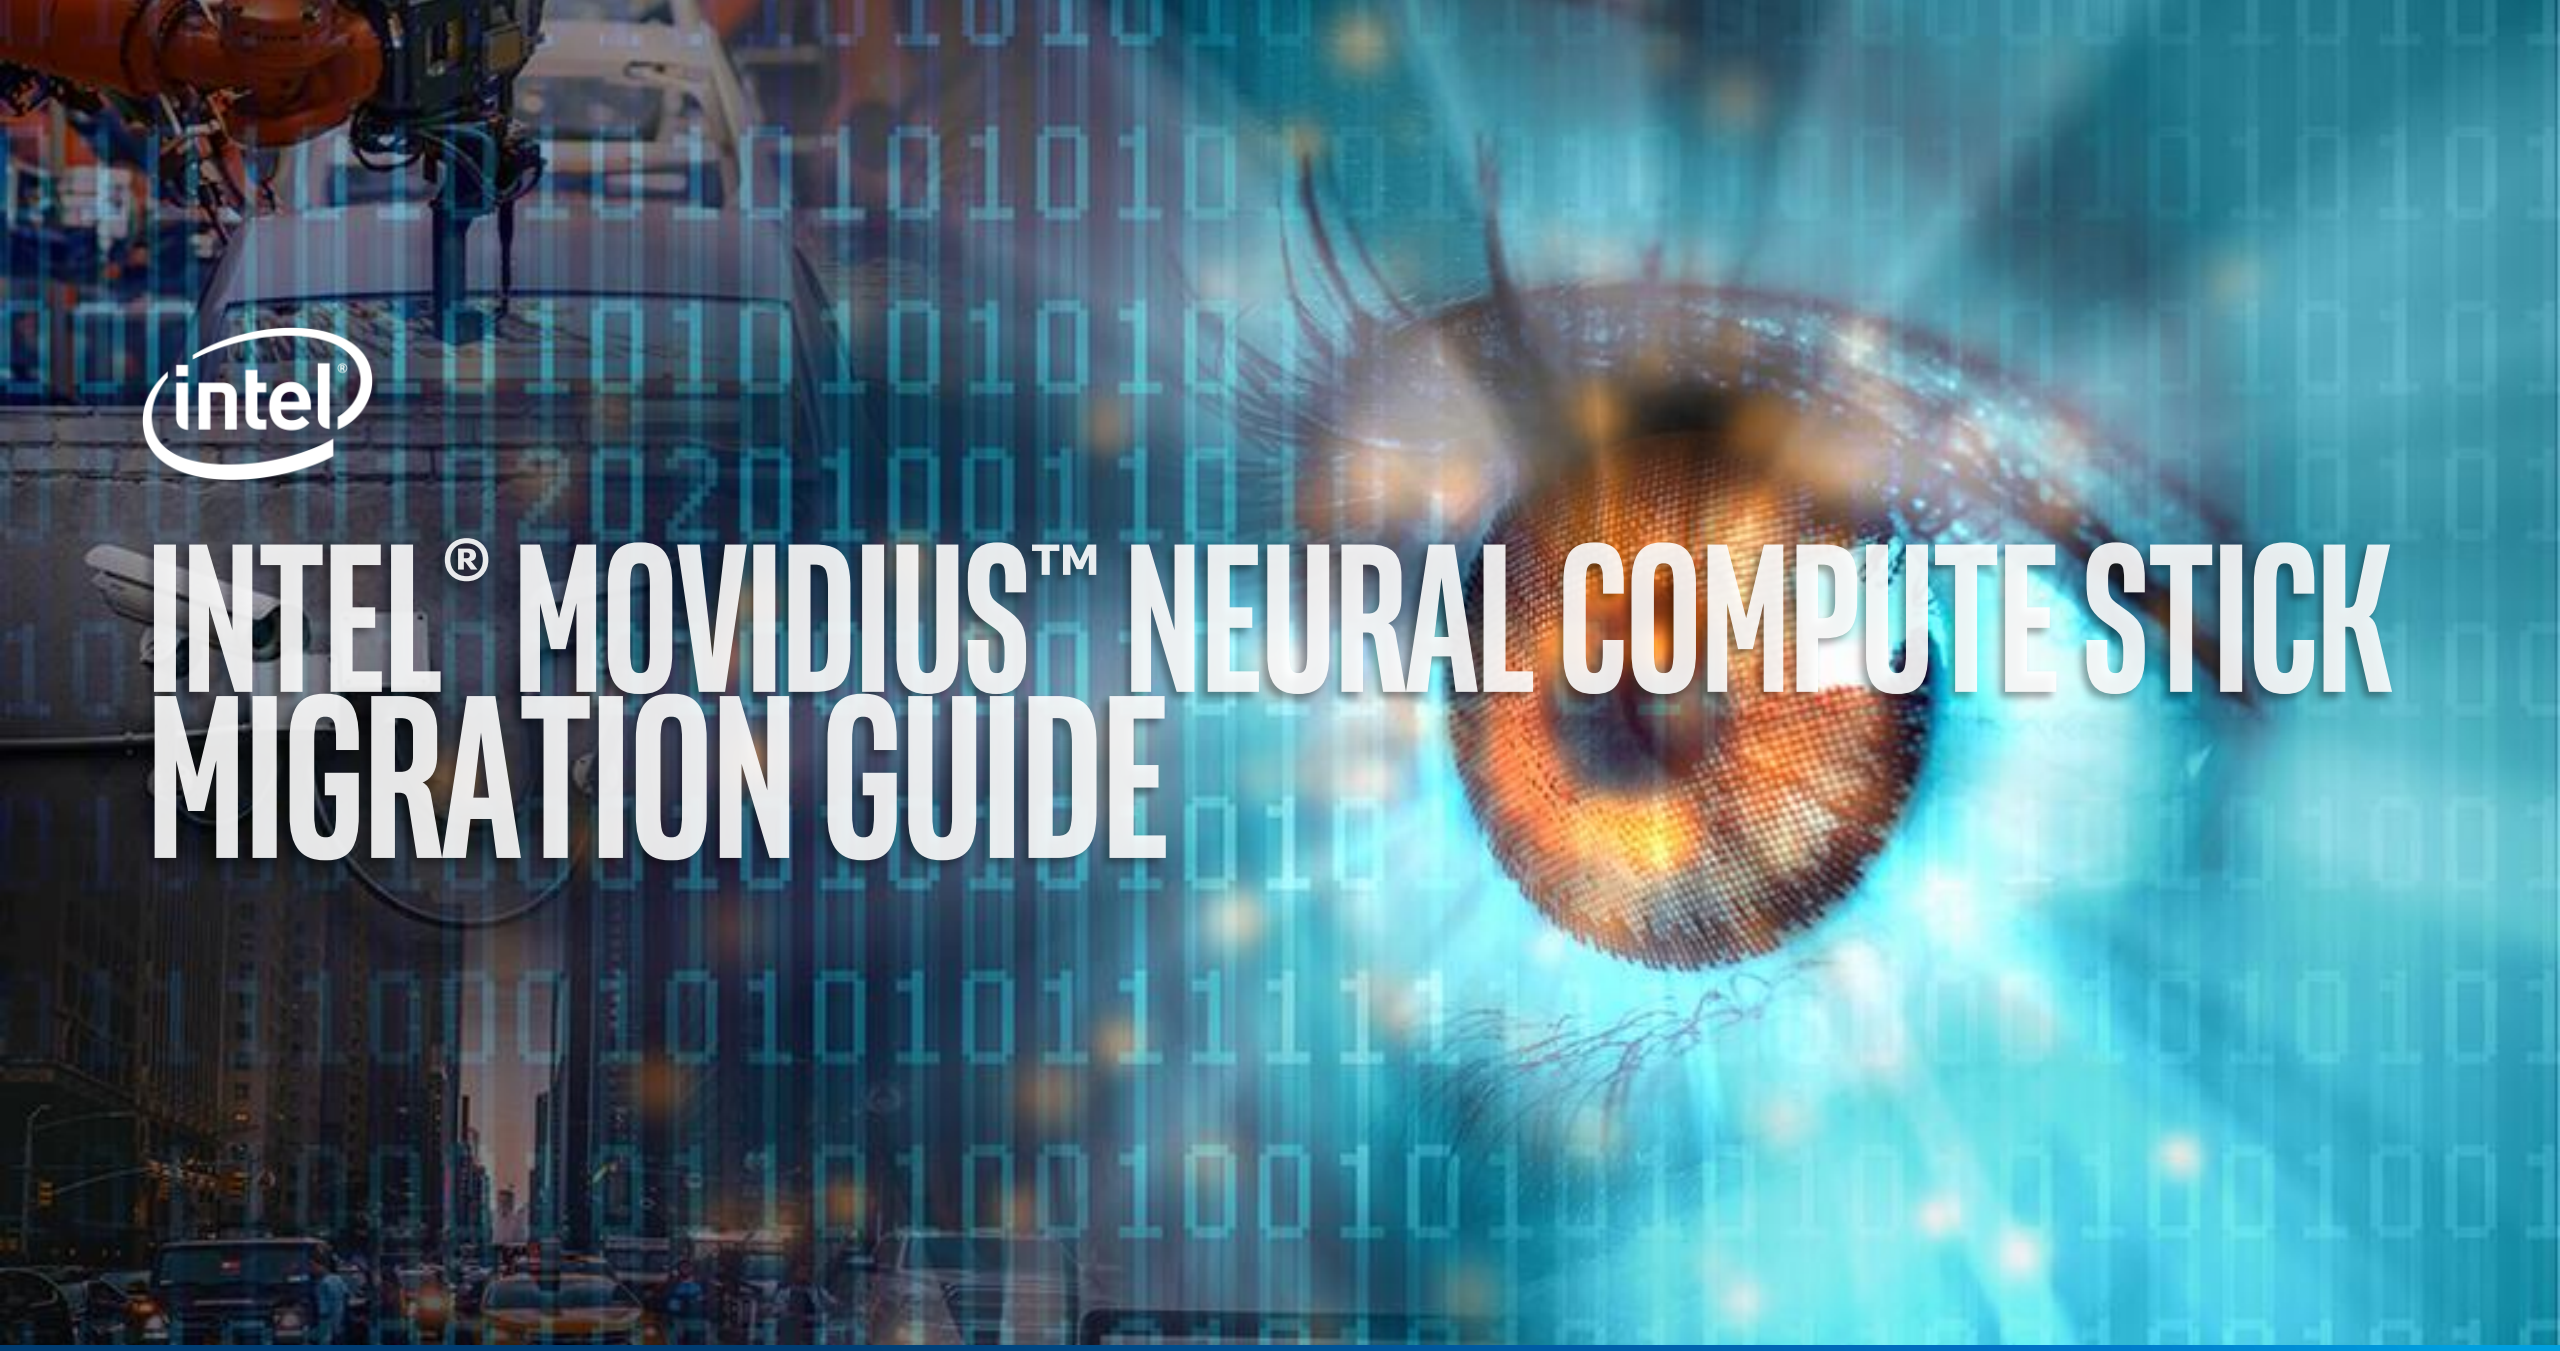

INTEL® MOVIDIUS™ NEURAL COMPUTE STICK MIGRATION GUIDE

INTEL® MOVIDIUS™ NEURAL COMPUTE STICK
PRODUCT CHANGE NOTIFICATION

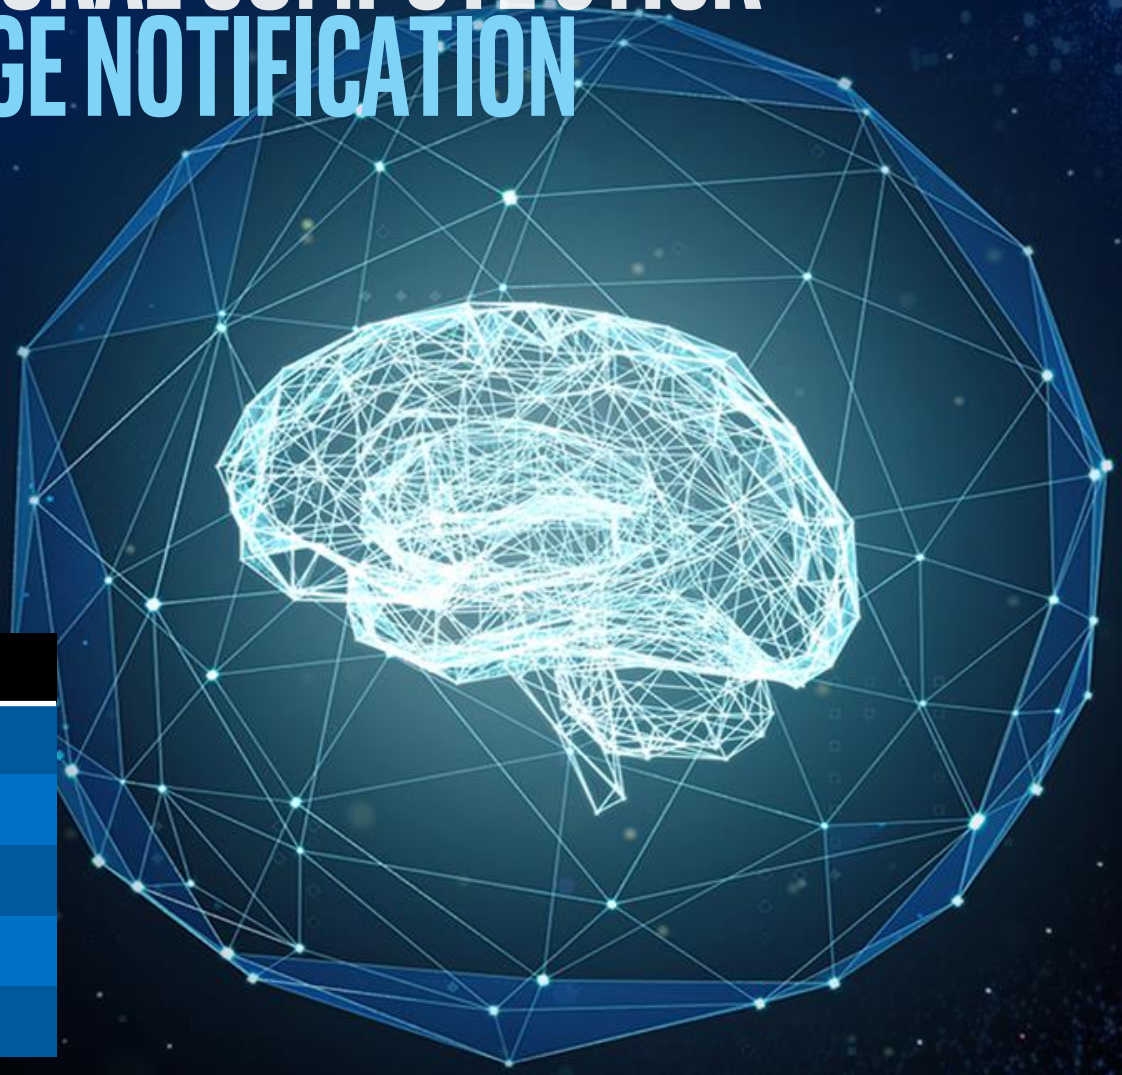

Intel will discontinue the Intel® Movidius™ Neural Compute Stick (MM# 962297).

Please contact your Intel Field Sales Representative with any questions, requests or concerns.

Key Milestones	Date
Product Discontinuance Program Support Begins	April 30, 2019
Last Product Discontinuance Order Date	October 30, 2019
Last Product Discontinuance Shipment Date	April 30, 2020
Technical Support Ends	April 30, 2021
Warranty Support Ends	April 30th, 2022

THE ORIGINAL INTEL® MOVIDIUS™ NEURAL COMPUTE STICK
UNLOCKED DEEP LEARNING OPPORTUNITIES

- Miniature deep learning development kit
- In a USB stick form-factor
- To prototype, tune, validate and deploy Deep neural networks
- Features the same Movidius vision processing unit (VPU) used in drones, surveillance cameras, VR headsets, and other low-power intelligent and autonomous products

NEXT-GENERATION AI AT THE EDGE REQUIRES SMARTER AND FASTER CAPABILITIES

45% *of data will be stored, analyzed, and acted on at the edge by 2019[1]*

43% *share of AI tasks taking place on edge devices (vs. cloud) in 2023[2]*

15X *growth in devices with edge AI capabilities by 2023[2]*

- **Increasing need for AI at the edge** requires exceptional performance per watt, hardware processing versatility and streamlined developer experience
- Now with enhanced hardware processing capabilities, the next generation **Intel® Neural Compute Stick 2** brings up to **8x[1] performance gain** on deep neural network inference
- **Intel® Distribution of OpenVINO™ toolkit** extends workloads across different hardware and maximizes performance

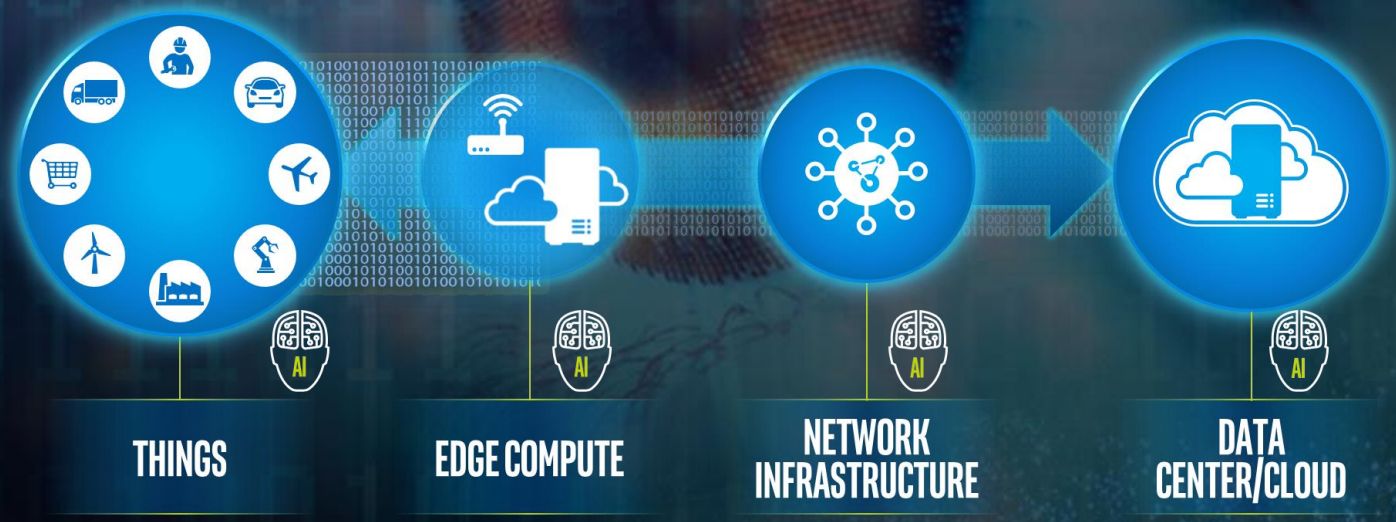

THE JOURNEY HAS JUST BEGUN
MIGRATION OPTIONS

INTEL® NEURAL COMPUTE STICK 2

Enhanced hardware processing capabilities vs. the original Intel® Movidius™ Neural Compute Stick; take advantage of **16 cores** instead of the original 12, plus a **neural compute engine**—a dedicated deep neural network accelerator

The new SDK, **Intel® Distribution of OpenVINO™ toolkit**, streamlines the development experience

Buy Now ▶

VISION ACCELERATOR KITS

IEI TANK AIoT Development Kit

Pre-installed with the Intel® Distribution of OpenVINO™ toolkit with great flexibility for hardware expansion on Intel® Arria™ 10 FPGA

AAEON UP Squared AI Vision Developer Kit

Pre-installed with the Intel® Distribution of OpenVINO™ toolkit with great flexibility for hardware expansion on Intel® Movidius™ Myriad™ X VPU

INTEL® VISION ACCELERATOR DESIGN

Intel® Movidius™ Myriad™ X VPU plug-in

High compute efficiency, low power and form factor constraints (e.g., cameras), and excellent performance/W/$, for well-defined workloads

Intel® Arria™ 10 FPGA plug-in

TOPS performance required on a single chip, support compute intensive networks (VGG*, ResNet* 101), and ensure continual performance optimization

Explore Products ▶

1 CONTINUE PROTOTYPING

2 PRODUCTIZE

3 INCREASE PERFORMANCE

HIGH PERFORMANCE & LOW POWER FOR AI INFERENCE
INTEL® NEURAL COMPUTE STICK 2

BOOST PRODUCTIVITY

SIMPLIFY PROTOTYPING

DISCOVER EFFICIENCIES

Powered by

Intel® Movidius™ Myriad™ X VPU

+

Optimized by

OpenVINO™

Intel® Distribution of OpenVINO™ toolkit

UP TO 8X[1] HIGHER PERFORMANCE

On deep neural networks compared to Intel® Movidius™ Neural Compute Stick

MORE CORES. MORE AI INFERENCE.

- ✓ Start quickly with plug-and-play simplicity
- ✓ Develop on common frameworks and out-of-box sample applications
- ✓ Prototype on any platform with a USB port
- ✓ Operate without cloud compute dependence

Order now for $99 MSRP*: Where to buy

*MSRP is not a guarantee of final retail price. MSRP may be changed in the future based upon economic conditions.

HIGH PERFORMANCE & LOW POWER FOR AI INFERENCE
INTEL® NEURAL COMPUTE STICK 2

Specifications	Intel® Movidius™ Neural Compute Stick	Intel® Neural Compute Stick 2
Vision Processing Unit (VPU)	The Intel® Movidius™ Myriad™ 2 VPU	The Intel® Movidius™ Myriad™ X VPU
Software development kit	The Intel® Movidius™ Neural Compute SDK	The Intel® Distribution of OpenVINO™ toolkit
OS support	Ubuntu* 16.04, Raspberry Pi* 3 Model B running Stretch desktop or Ubuntu 16.04 Virtual Box instance	Ubuntu* 16.04.3 LTS (64 bit), Windows 10 (64 bit), or CentOS 7.4 (64 bit)
Supported framework	TensorFlow* and Caffe*	TensorFlow* and Caffe*
Connectivity	USB 3.0 Type-A	USB 3.0 Type-A
USB stick dimensions (mm)	72.5mm X 27mm X 14mm	72.5mm X 27mm X 14mm
Operating temperature	0° - 40° C	0° - 40° C
Material Master Number	962297	964486
MSRP	$79 USD*	$99 USD*

*MSRP is not a guarantee of final retail price. MSRP may be changed in the future based upon economic conditions.

INTEL® DISTRIBUTION OF OPENVINO™ TOOLKIT

ENABLES WIDE DEPLOYMENT OF DEEP LEARNING ALGORITHMS

DEEP LEARNING

Caffe — >100 Networks Supported

mxnet — CV Algorithms

ONNX TensorFlow — Model Optimizer

Inference Engine

COMPUTER VISION

OpenCV

OpenCL

OpenVX™

CV Library (Kernel & Graphic APIs)

AGNOSTIC, COMPLEMENTARY TO MAJOR FRAMEWORKS | CROSS-PLATFORM FLEXIBILITY | HIGH PERFORMANCE, HIGH EFFICIENCY

Over 6500 unique developers

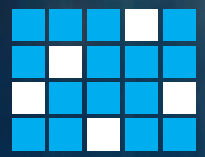

Over 20 customer products launched based on OpenVINO™ toolkit

Breadth of vision product portfolio, and now, introducing...

STRONG ADOPTION + RAPIDLY EXPANDING CAPABILITY

TRANSITION TO INTEL® DISTRIBUTION OF OPENVINO™ TOOLKIT
FROM INTEL® MOVIDIUS™ NEURAL COMPUTE SDK

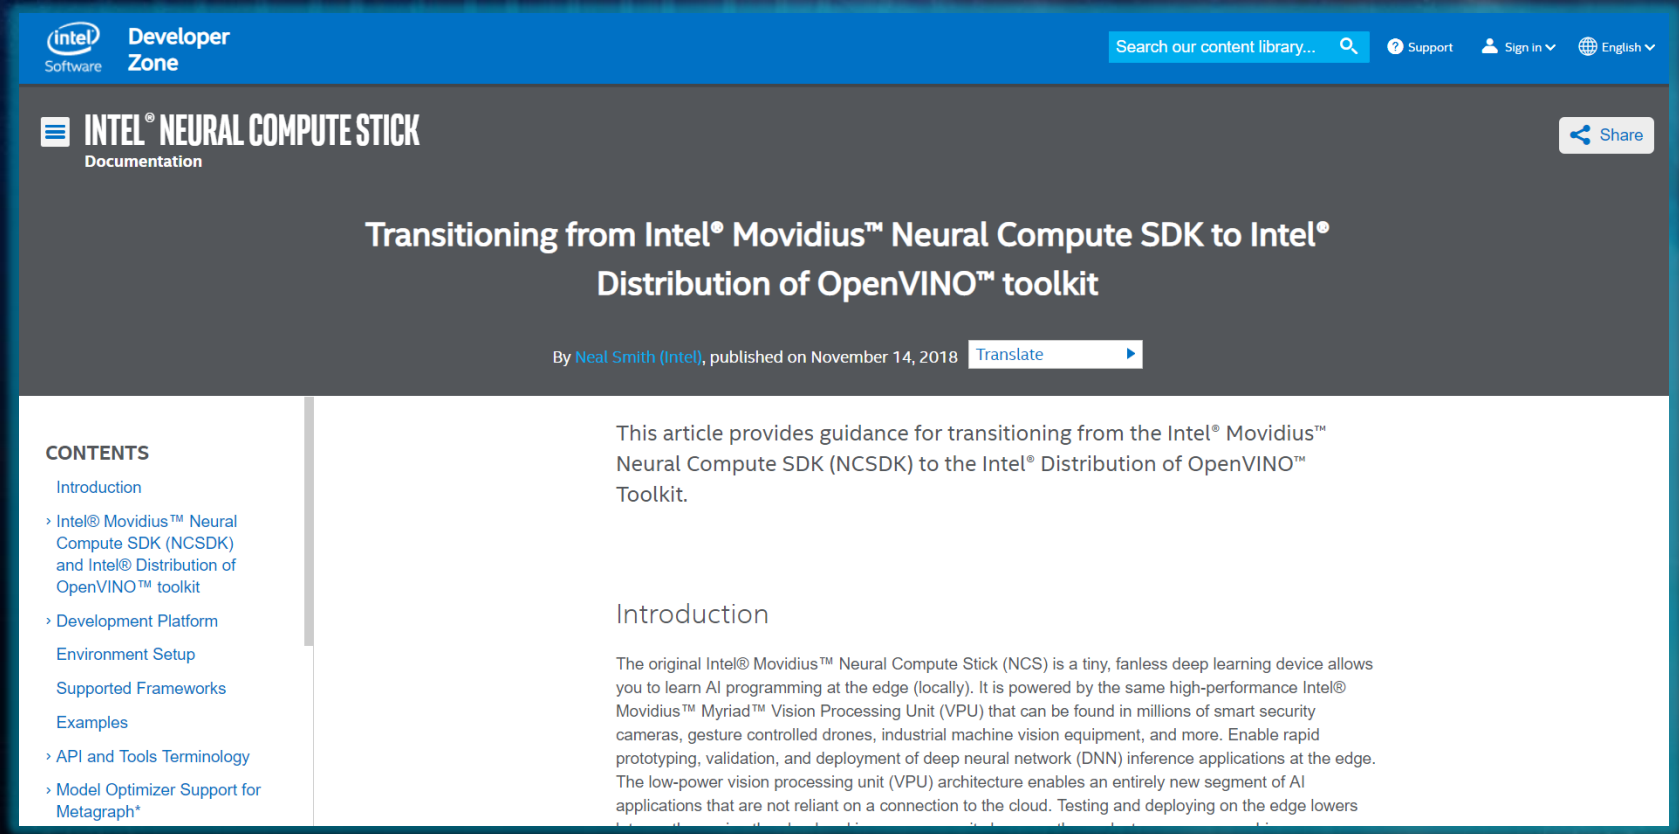

[Transition Guide to Intel® Distribution of OpenVINO™ toolkit]

VISION ACCELERATOR KITS

REDUCE TIME TO PROTOTYPE | EXPEDITE PATH TO PRODUCTION | FUTURE PROOF DESIGN

UP SQUARED AI VISION
Development Kit

bridge the gap

- UP Squared Board (Intel® Atom x7-E3950)
- AI Core- Myriad X mini PCIe
- USB camera
- **Intel® Distribution of OpenVINO™ toolkit**
- Ubuntu 16.04 Desktop w/ 4.14 kernel
- Intel® System Studio
- Arduino Create*

http://www.up-board.org/upkits/up-squared-ai-vision-kit/

IEI TANK AIoT
Developer Kit

TANK
ACCELERATE TO THE FUTURE

- IEI TANK 870-Q170
  (6th/7th Gen Intel® Core™ processor)
- **Intel® Distribution of OpenVINO™ toolkit**
- Intel® Media SDK
- Ubuntu Desktop16.04 Desktop LTS
- Intel® System Studio

https://www.ieiworld.com/tank-aiot-development-kit/en/

Article on Intel Developer Zone: https://software.intel.com/en-us/kits-to-accelerate-your-computer-vision-deployments

INTEL® VISION ACCELERATOR DESIGN PRODUCTS

HOST IA PLATFORMS: APPLICATION PROCESSING, MEDIA, "FREE" CV/DL

Use the Intel® Media SDK to achieve en/de/trans-code performance

Maximize CV/DL performance on the host platform with the Open Visual Inference & Neural Network Optimization (OpenVINO™) toolkit

INTEL® MOVIDIUS™ VPUS

OVERVIEW

Intel Movidius VPUs offer high performance per watt per dollar.
Easily add AI-based visual intelligence by plugging in one or more cards.

VALUE PROP

Intel Movidius VPUs enable deep neural network inferencing workloads with high compute efficiency, low power and form factor constraints (e.g., cameras), and excellent performance/W/$, for well-defined workloads.

KEY USE CASES

Intel Movidius VPUs work well with networks that have:

- A small memory footprint (less than 250 MParameters)
- Lower performance requirements (<3 GMACs)
- FP16 precision
- Accelerator Power Budget: 2-25W
- # of streams: 1-16

INTEL® ARRIA® 10 FPGAS

OVERVIEW

Intel Arria 10 FPGAs offer exceptional performance, flexibility, and scalability for NVRs, edge deep learning inference appliances, and on-premise servers or cloud.

VALUE PROP

Intel Arria 10 FPGAs achieve TOPS performance required on a single chip, support compute intensive networks (VGG*, ResNet* 101), and ensure continual performance optimization by taking advantage of quarterly bitstream updates from Intel to reduce the need to upgrade hardware.

KEY USE CASES

The Intel Arria 10 FPGAs work well with networks that have:

- Larger memory footprint (more than 250 MParameters)
- Larger performance requirements (>3 GMACs)
- FP16/11/9 precision
- Accelerator Power Budget: <50W
- # of streams: 3-15

INTEL® VISION ACCELERATOR DESIGN PRODUCTS

AI AT THE EDGE ACCELERATION

Equipment Maker	Product Name	Form Factor	Platform	Board Dimensions	Availability
AAEON Technology Inc.	AI Core X	mPCIe	1x Intel® Movidius™ Myriad™ X 2485 VPU	51 x 30 mm	Now
	AI Core XM 2280	M.2 2280 B+M key	2x Intel® Movidius™ Myriad™ X 2485 VPUs	22 x 80 mm	Now
	AI Vision Plus X	custom	3x Intel® Movidius™ Myriad™ X 2485 VPUs	90 x 56.5 mm	Q2'19
ADLINK	EDL-mPCIe-MA2485	mPCIe	2x Intel® Movidius™ Myriad™ X 2485 VPUs	30 x 50 mm	Q2'19
IEI Integration Corp.	VEGA-320	M.2 2230 B+M key	1x Intel® Movidius™ Myriad™ X 2485 VPU	22 x 30 mm	Now
	VEGA-330	mPCIe	2x Intel® Movidius™ Myriad™ X 2485 VPUs	22 x 30 mm	Now
	Mustang-V100-MX8-R10	PCIe x4	8x Intel® Movidius™ Myriad™ X 2485 VPUs	169.54 x 56.16 mm	Now
	Mustang-F100-A10-R10	PCIe x8	1x Intel® Arria® 10 FPGA 1150 GX	169.54 x 58.40 mm	Now
Nexcom	AIBooster-X8	PCIe x4	8x Intel® Movidius™ Myriad™ X 2485 VPUs	169.54 x 58.40 mm	Q2'19
Shenzhen Uzel Information Technology Co., Ltd.	HDDL-L	mPCIe	2x Intel® Movidius™ Myriad™ X 2485 VPUs	28.00*50.90 mm	Now
	HDDL-R	PCIE x4	8x Intel® Movidius™ Myriad™ X 2485 VPUs	68.90*157.00 mm	Now
	HDDL-S	Mother board	24x Intel® Movidius™ Myriad™ X 2485 VPUs and 1*intel Core I3/I5/I7/E3 CPU	135.00*445.00 mm	Q2'19

Back to Overview Page

RESOURCES

- Intel® Neural Compute Stick 2 website
- Intel® Neural Compute Stick 2 tutorials
- Where to purchase the Intel® Neural Compute Stick 2
- Intel® Neural Compute Stick 2 Get Started page
- Intel® Distribution of OpenVINO™ toolkit homepage
- Intel® Distribution of OpenVINO™ toolkit Pre-trained Models
- Reference Implementations

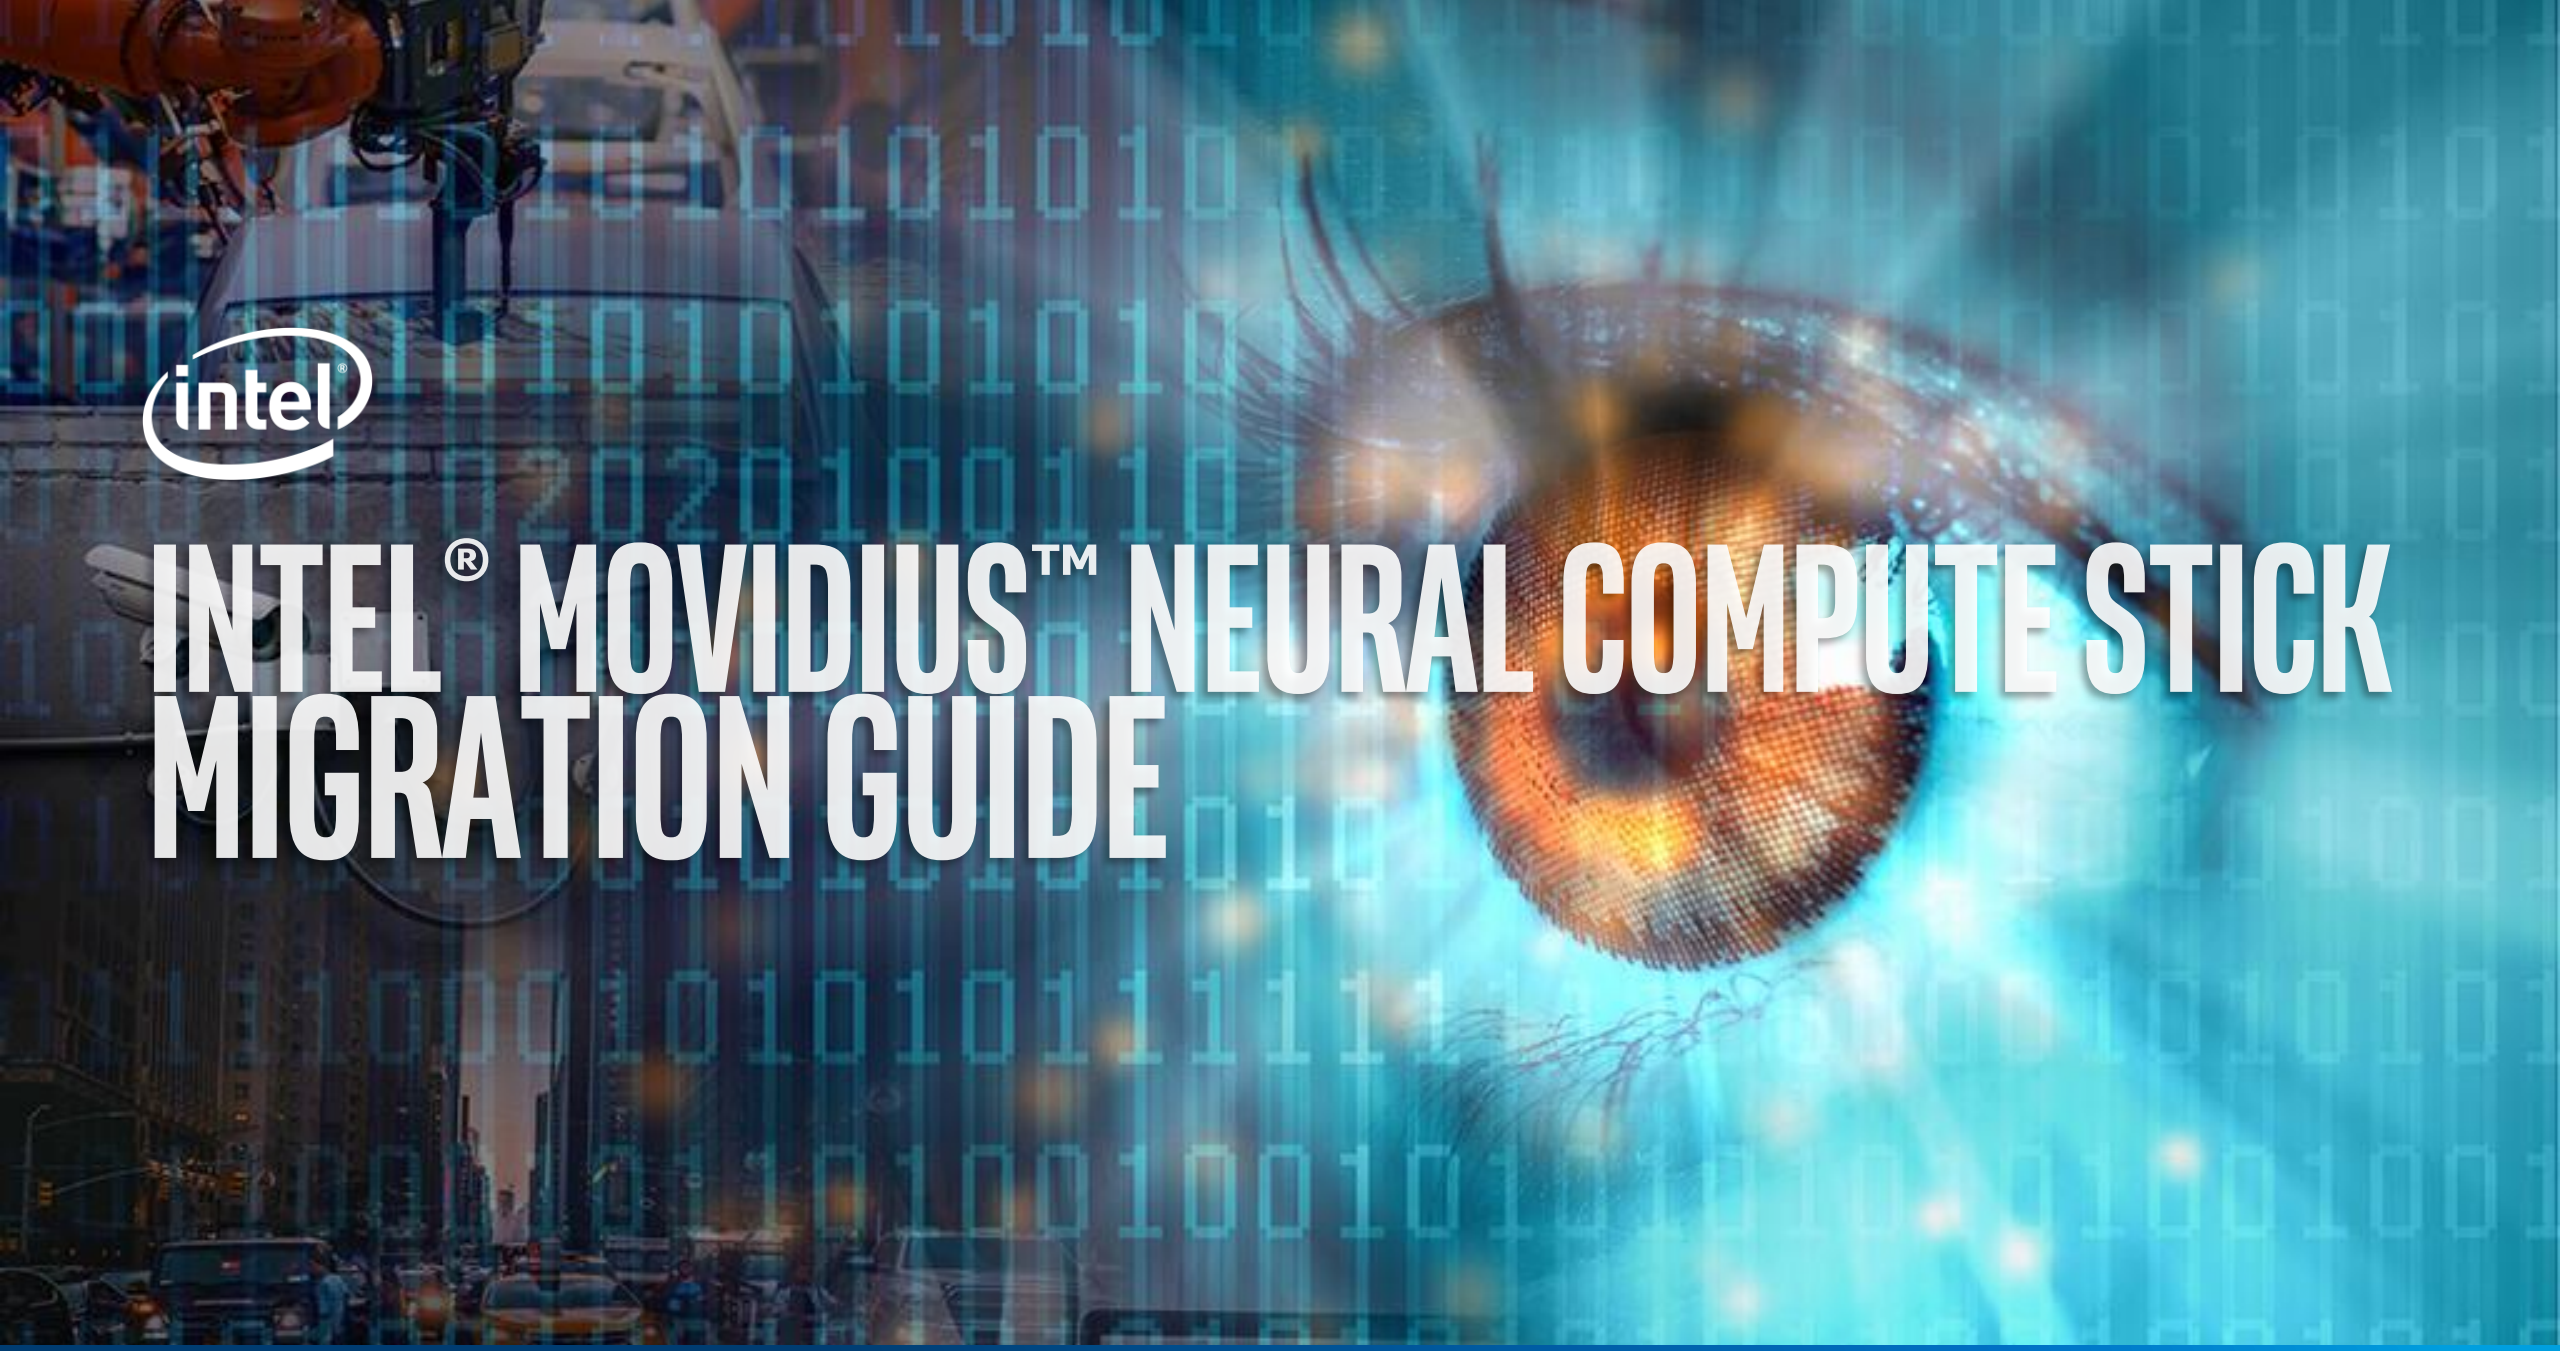

INTEL® MOVIDIUS™ NEURAL COMPUTE STICK MIGRATION GUIDE

LEGAL DISCLAIMER & OPTIMIZATION NOTICE

[1]Testing by Intel as of October 12th, 2018

Deep Learning Workload Configuration. Comparing Intel® Movidius™ Neural Compute Stick based on Intel® Movidius™ Myriad™ 2 VPU vs. Intel® Neural Compute Stick 2 based on the Intel® Movidius™ Myriad™ X VPU with Asynchronous Plug-in enabled for (2xNCE engines). As measured by images per second across GoogleNetV1*. Base System Configuration: Intel® Core™ I7-8700K 95W TDP (6C12T at 3.7GHz base freq and 4.7GHz max turbo freq), Graphics: Intel® UHD Graphics 630 Total Memory 65830088 kB Storage: INTEL SSDSC2BB24 (240GB), Ubuntu* 16.04.5 Linux*-4.15.0-36-generic-x86_64-with-Ubuntu*-16.04-xenial, deeplearning_deploymenttoolkit_2018.0.14348.0, API version 1.2, Build 14348, myriadPlugin, FP16, Batch Size = 1

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 visit www.intel.com/benchmarks.

Performance results are based on testing as of October 12th, 2018 and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure.

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.

The benchmark results reported in this deck may need to be revised as additional testing is conducted. The results depend on the specific platform configurations and workloads utilized in the testing, and may not be applicable to any particular user's components, computer system or workloads. The results are not necessarily representative of other benchmarks and other benchmark results may show greater or lesser impact from mitigations.

INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS". NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO THIS INFORMATION INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

OpenVINO and the OpenVINO logo 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.

Optimization Notice

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