
GET SOFTWARE SMART!
Olga Andrianova
SOFTWARE BUSINESS GROWTH FOR
RUSSIA/EASTERN EUROPE COUNTRIES


FOR YOUR BUSINESS. BY DEVELOPERS.
• Are You Ready to Reinvent Your Software Solution?
Best performance in AI, Gaming, Big data and analytics, Fast code, HPC&Cloud, Security, Networking, Virtual Reality, IoT, embedded, Open source
• Through industries to end user experience!
Software - hardware, interoperability, and amazing customer experiences
• Intel® software is your partner !
Easy , open and scalable solution on Intel systems


COMMITMENT TO TECHNOLOGY ADVANCEMENT
Learning and Teaching!
Workshops, deep dive trainings, students schools
Industry collaboration - join us to and discover how software + hardware work together !
End-to end solutions with our partners!

Olga.Andrianova@intel.com

+7 903 042 11 25

https://software.intel.com/

INTEL: THE FUTURE OF TECHNOLOGY: AUTONOMOUS DRIVING, AI, BIG DATA

Ivan Kuzmin

DIRECTOR OF INTEL PERFORMANCE LIBRARIES GROUP

RUSSIA

AI IS EXPLODING

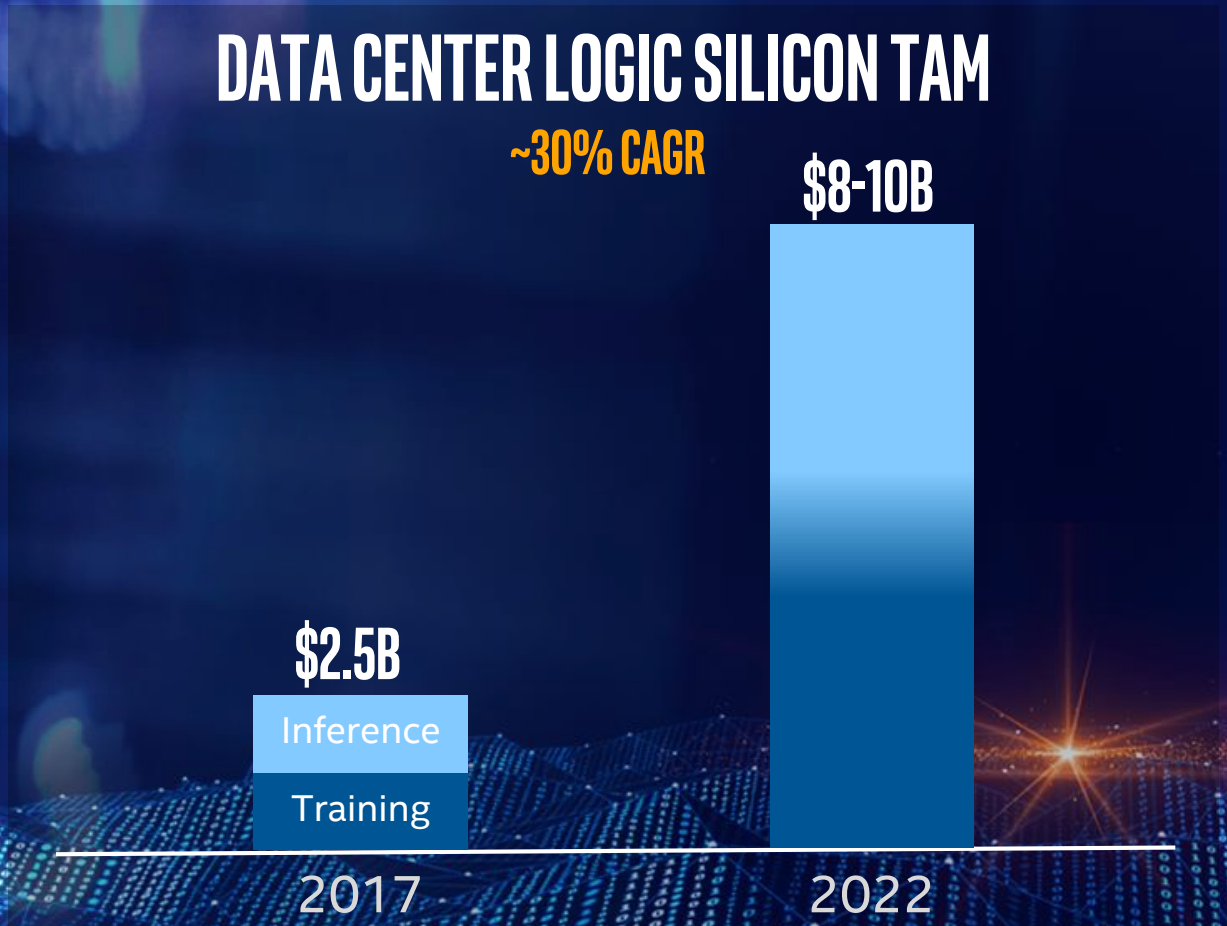

EMERGING AS A CRITICAL WORKLOAD

1. Source: AI Si Server TAM is based on amalgamation of analyst data and Intel analysis, based upon current expectations and available information and are subject to change without notice.

AI IS EVOLVING

PROOFS OF CONCEPTS → UNLOCKING REAL VALUE

AI IS EXPANDING

END POINT | EDGE | DATA CENTER

intel MOVIDIUS inside | MobilEye An Intel Company | intel ATOM inside | intel FPGA inside | intel CORE i7 inside | intel NERVANA inside | intel XEON PLATINUM inside

COMPREHENSIVE AI PORTFOLIO

AI EXAMPLES

Variety of AI approaches for different problem types

EXAMPLES

DEEP LEARNING

Image/speech recognition, natural language processing, pattern recognition/detection , etc.

CLASSIC ML

Statistical problems, recommendation engines, transparency requirement, etc.

REASONING

Multivariate supply chain probe, full database fraud detection, whole CRM churn analysis, etc.

EMERGING

AI research: ‘sequence alignment’ in computational biology, ‘binary neural network based inferencing’, etc.

END-TO-END EXAMPLE

Automated Driving

VEHICLE

Driving Functions
Autonomous Driving Functions
Trajectory Enumeration, Path Planning, Selection & Maneuvering
Driving Policy, Path Selection

Real Time Environment Modeling
Localization

Sensor Processing & Fusion
Object ID & Classification

Anomaly Detection

NETWORK

Compressed Models

OTA SW /FW Updates

Data Formatting

DATACENTER

Data Storage
Dataset Management
&
Traceability

Neural Network Design for Target Hardware, & usage
(Vision, Data Driven, etc.)

Model Training
Single & Multi-node optimized Frameworks

Model Inference >than Real Time
Model Simulation & Verification

ONE SIZE DOES NOT FIT ALL

END POINT

IOT SENSORS
(Security, home, retail, industrial...)

Intel Movidius inside — Vision & Inference

Intel Atom inside + Intel GNA inside — Speech

SELF-DRIVING VEHICLE

Mobileye — Autonomous Driving

DESKTOP & MOBILE

Intel Core i7 inside + Intel Iris + Intel GNA inside + Intel Movidius inside

Display, video, AR/VR, gestures

CONVERGED MOBILITY

Intel Atom inside + Intel Movidius inside

Vision, speech, AR/VR

EDGE

SERVERS, APPLIANCES & GATEWAYS

Intel Xeon Platinum inside — Most use cases

Intel Arria 10 inside — Streaming latency-bound systems

Intel Movidius inside — Vision & Inference for various systems types

DATA CENTER

SERVERS & APPLIANCES

Intel Xeon Platinum inside — Foundation for AI

Intel Nervana inside — Built for Deep Learning

Intel Stratix 10 inside — Flexible & memory bandwidth bound use cases

WINNING TOGETHER WITH INTEL AI

Subset of full customer and partner list

$1B+ AI BUSINESS FOR INTEL TODAY

intel

AI DEVELOPMENT LIFECYCLE

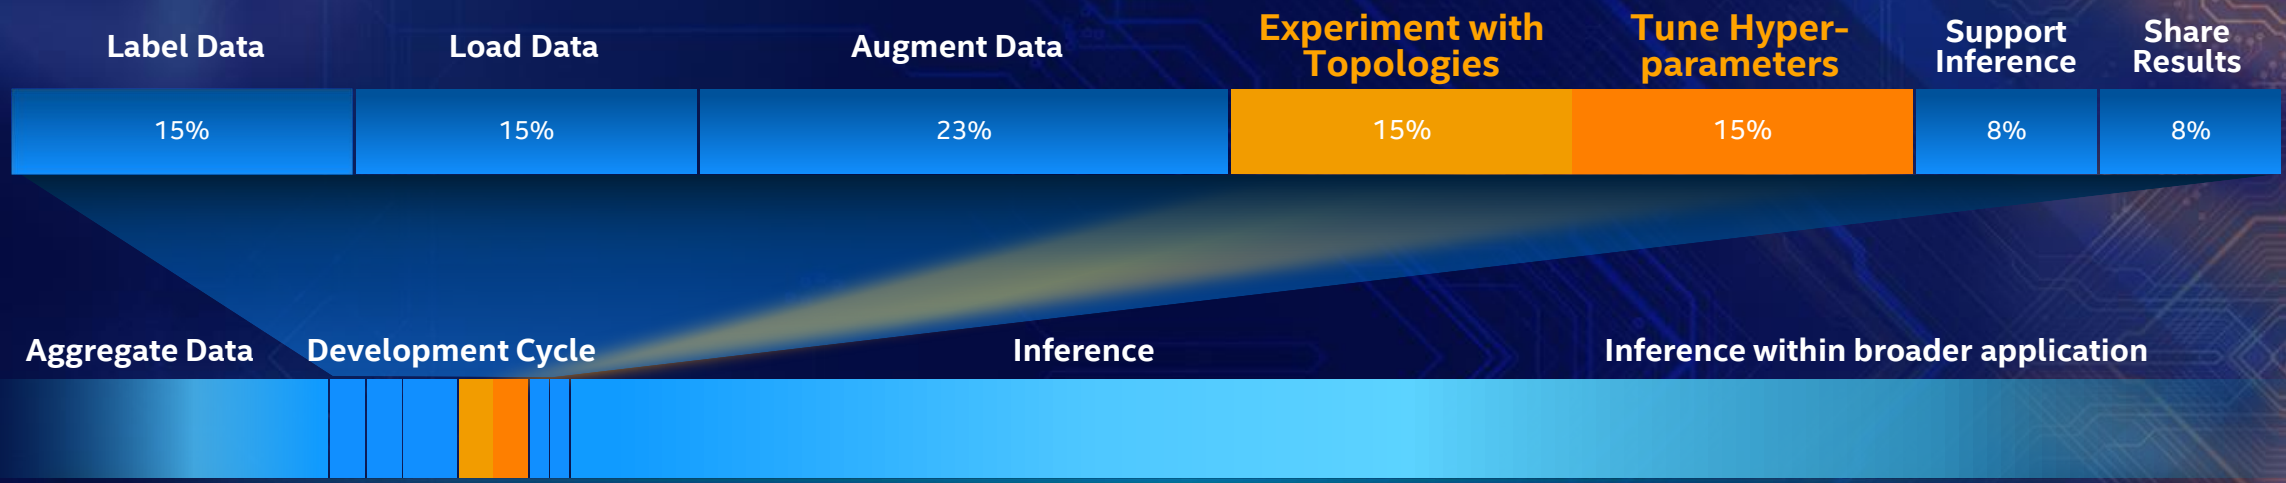

INTEL WORKS WITH CUSTOMERS ACROSS ENTIRE PROCESS FLOW

BROUGHT TO LIFE THROUGH DATA SCIENTISTS

RESEARCH CUSTOMIZE DEPLOY

INTEL® XEON® SCALABLE PROCESSORS

THE FOUNDATION FOR AI

INFERENCE

Projected Performance Intel® DL Boost with Vector Neural Network Instruction (VNNI)

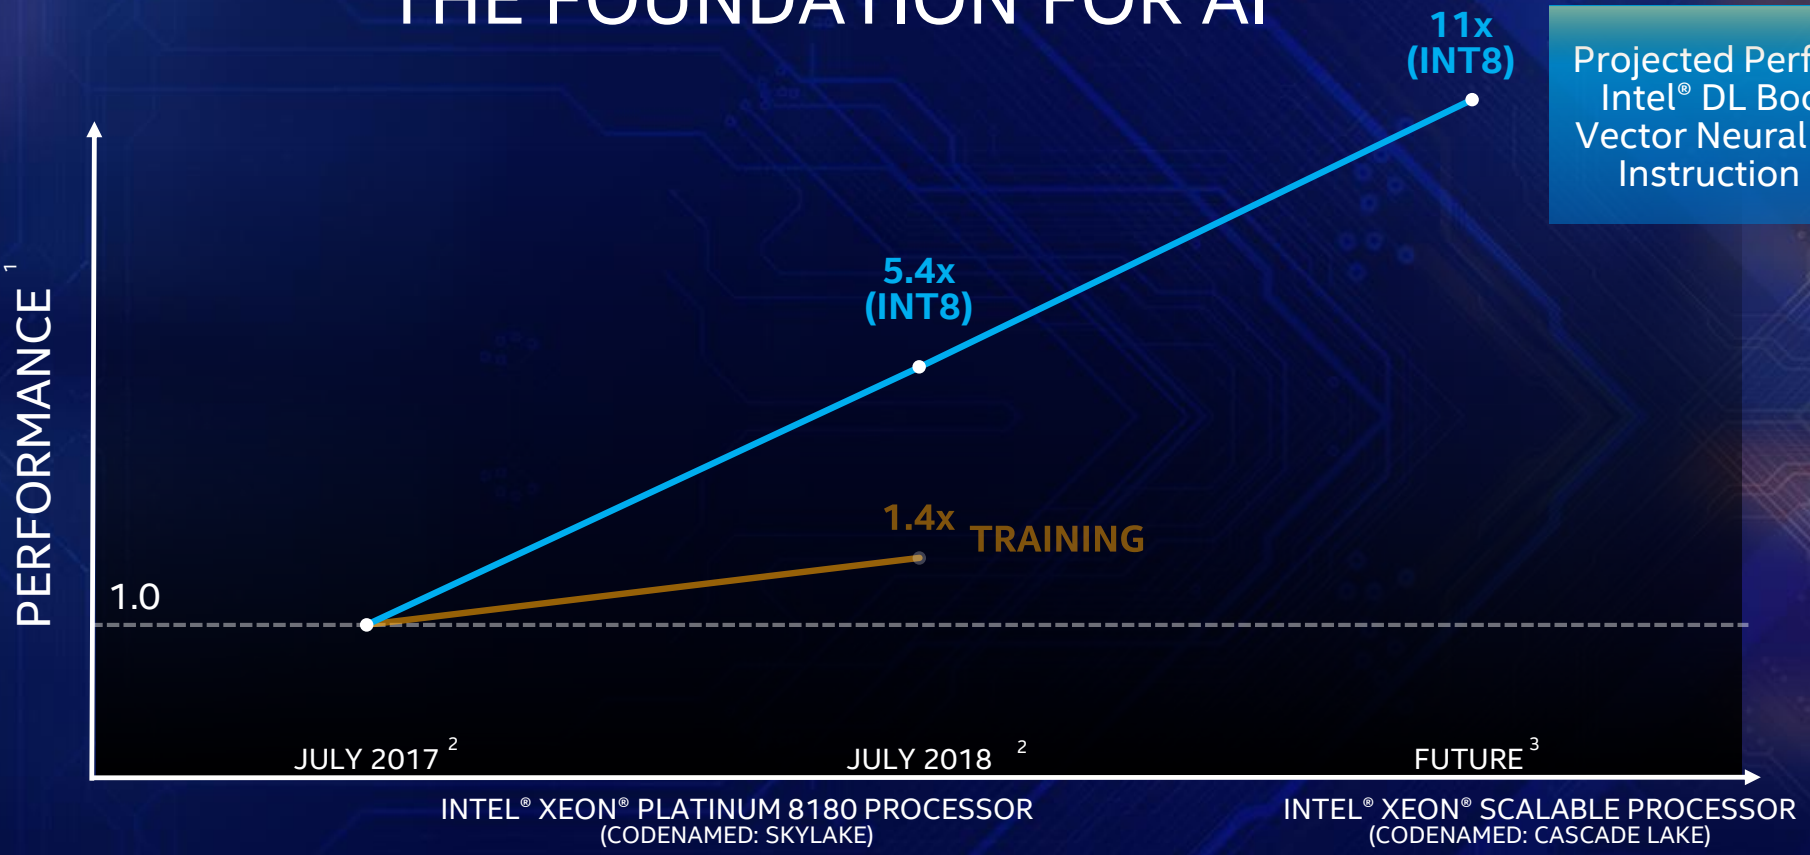

CONTINUED INVESTMENTS IN OPTIMIZATIONS TO DELIVER INCREASED PERFORMANCE

[1] Intel® Optimization for Caffe Resnet-50 performance does not necessarily represent other Framework performance.

[2] Based on Intel internal testing: 1X (7/11/2017), 2.8X (1/19/2018), 1.4x (8/2/2018) and 5.4X (7/26/2018) performance improvement based on Intel® Optimization for Café Resnet-50 inference throughput performance on Intel® Xeon® Scalable Processor. See Configuration Details Slide #19

Performance results are based on testing as of 7/11/2017(1x), 1/19/2018(2.8x), 8/2/2018 (1.4x) & 7/26/2018(5.4) and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure. 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.

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: http://www.intel.com/performance.

INTEL® NERVANA™ NNP L-1000

PURPOSE-BUILT FOR REAL WORLD AI PERFORMANCE

Optimized across memory, bandwidth, utilization and power

3-4x training performance of first-generation NNP product

High-bandwidth, low-latency interconnects

bfloat16 numerics

FIRST COMMERCIAL NNP IN 2019

Source: Based on Intel measurements on limited distribution SDV (codenamed: Lake Crest) compared to Intel measurements on NNP-100 simulated product

SOFTWARE IS ESSENTIAL

ABSTRACTION

TOOLKITS
Application Developers

OpenVINO™ Toolkit

Intel® Movidius™ SDK

LIBRARIES
Data Scientists

MACHINE LEARNING LIBRARIES

Scikit-Learn NumPy MLlib

DEEP LEARNING FRAMEWORKS

TensorFlow mxnet Caffe BigDL for Spark

Caffe2 PYTORCH Microsoft CNTK PaddlePaddle

FOUNDATION
Library Developers

ANALYTICS, MACHINE & DEEP LEARNING PRIMITIVES

MKL-DNN clDNN Python DAAL

DEEP LEARNING GRAPH COMPILER

Intel® nGraph™ Compiler

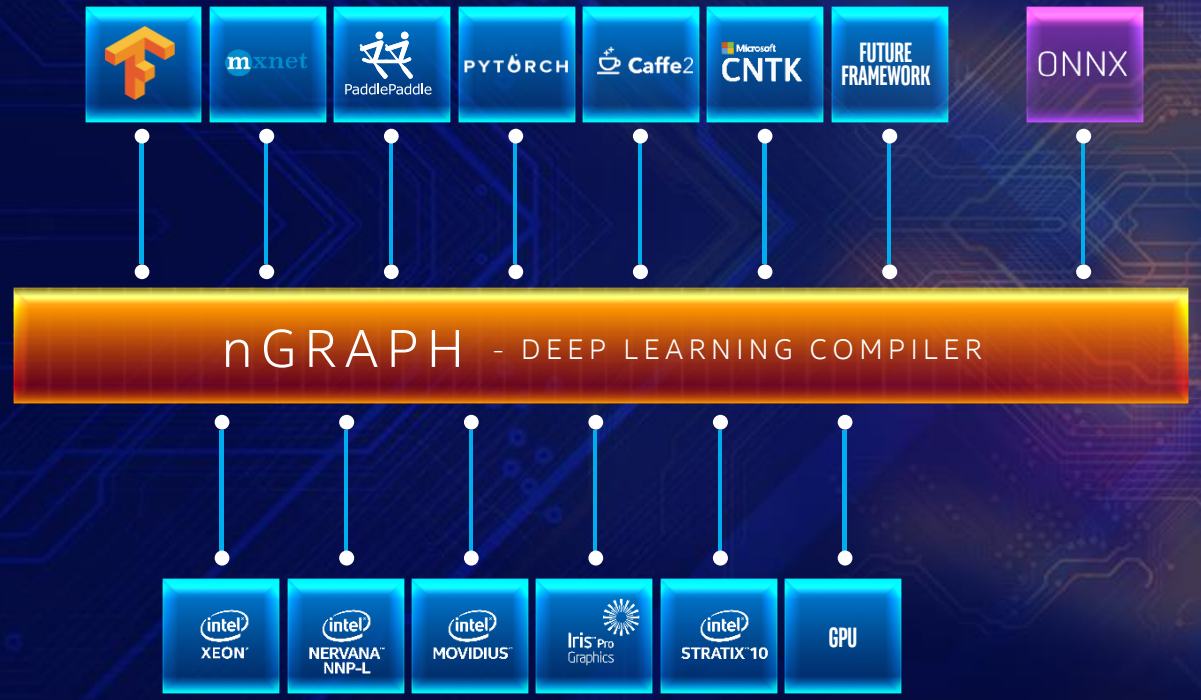

intel

BigDL FOR APACHE SPARK CLUSTERS

Problem: 90% of homebuyers search online with the top "wishlist" item being *"show me similar houses."*

Solution: Visual similarity neural network that helps home buyers find similar houses

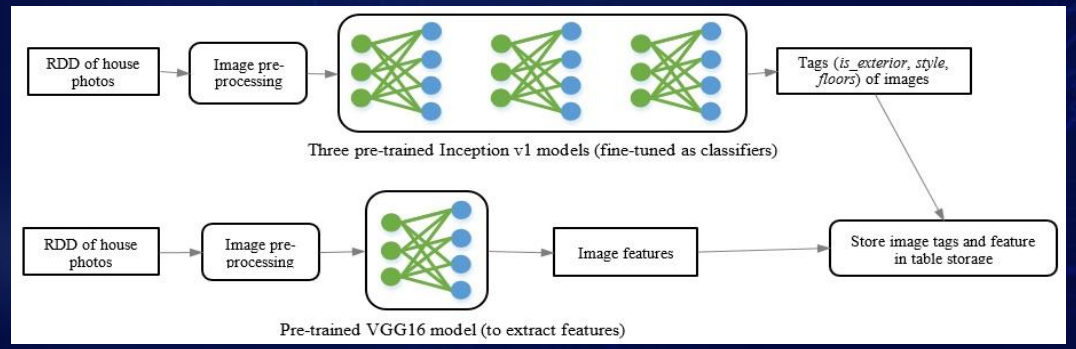

RESULTS AVAILABLE IN <0.1S

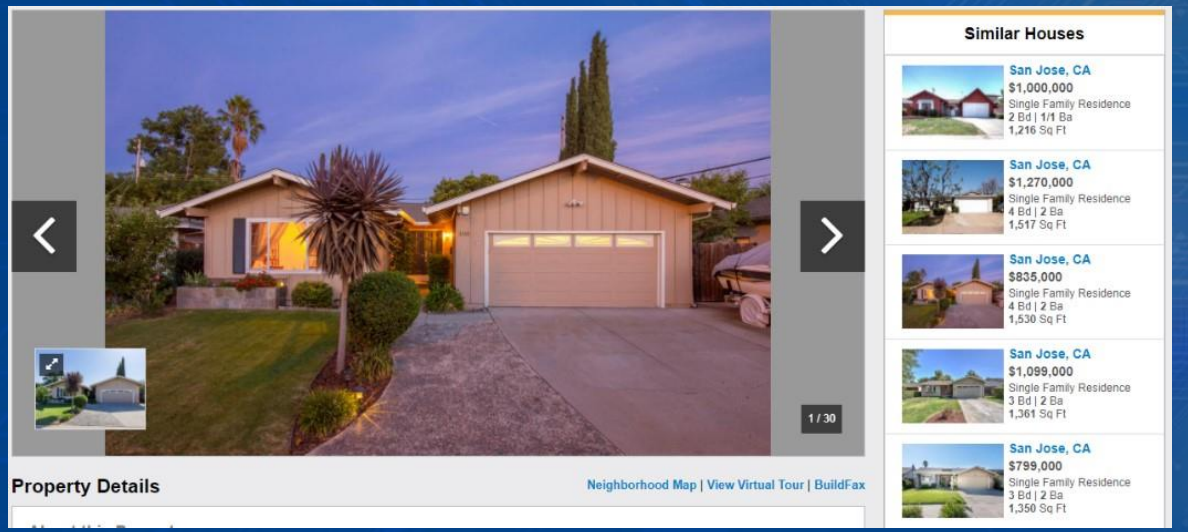

- Source: https://software.intel.com/en-us/articles/using-bigdl-to-build-image-similarity-based-house-recommendations.

OpenVINO™ Toolkit
FOR EDGE AND ENDPOINT DEPLOYMENTS

GE Healthcare

PROBLEM

- Need fast inferencing to classify images coming off of a CT scanner
- Improve workflow of the clinician, providing accurate images required for the examination

SOLUTION

- Low cost solution using existing Xeon® infrastructure
- High performance inferencing at ~596 images/sec
- Exceeded customer goals by 5.9x and provided 14x performance boost over unoptimized version*

CT SCAN IMAGES

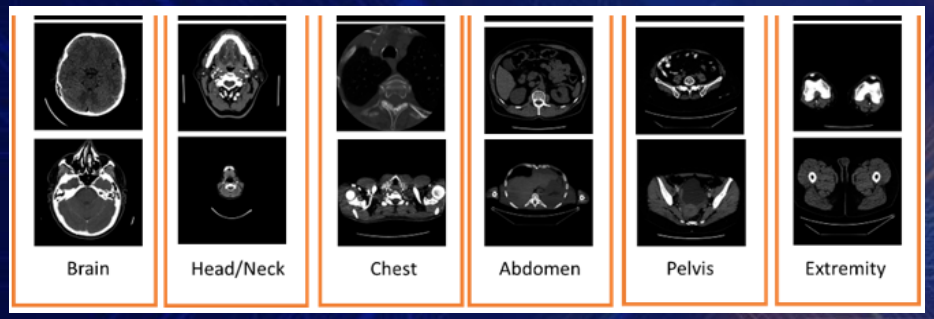

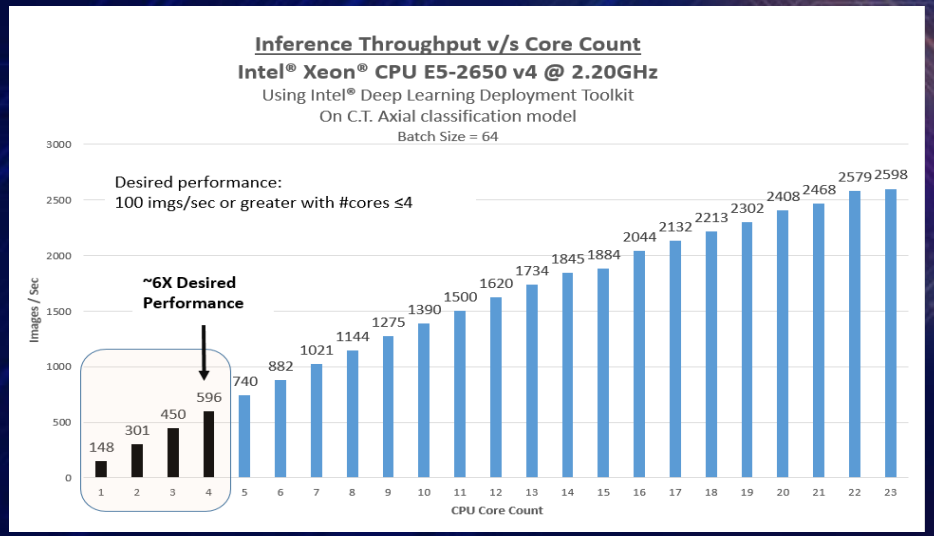

http://simplecore-dev.intel.com/ai/wp-content/uploads/sites/71/Intel-Software-Development-Tools-Optimize-Deep-Learning-Performance-for-Healthcare-Imaging-paper-ForDistribution.pdf

intel

HIGH PERFORMANCE AT SCALE

SCALING OF TIME TO TRAIN

INTEL® OMNI-PATH ARCHITECTURE, HOROVOD AND TENSORFLOW®

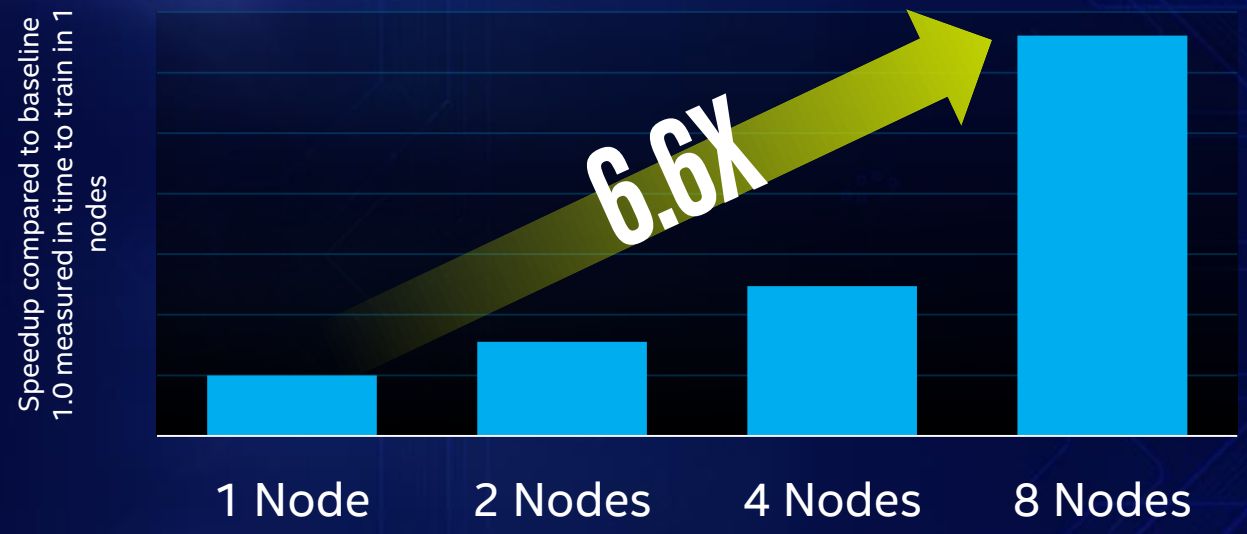

TOTAL MEMORY USED

192GB DDR4 PER INTEL® SP 2S XEON® 6148 PROCESSOR

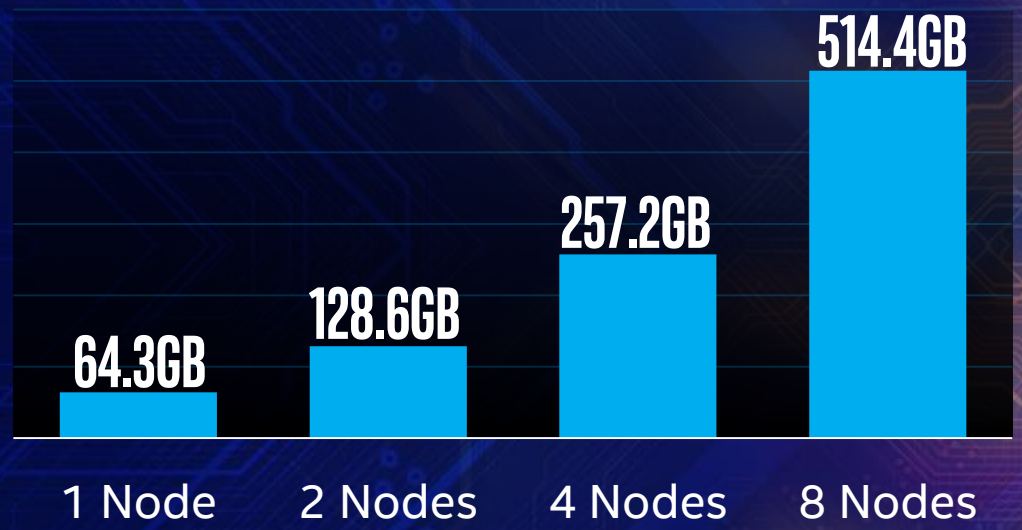

MULTISCALE CONVOLUTION NEURAL NETWORK

TensorFlow

OPTIMIZED LIBRARIES

Intel® MKL/MKL-DNN, clDNN, DAAL

INTEL® OMNI-PATH ARCHITECTURE

DATA-CENTRIC INNOVATION SUMMIT

§ Configuration: CPU: Intel Xeon 6148 processor @ 2.4GHz, Hyper-threading: Enabled. NIC: Intel® Omni-Path Host Fabric Interface, TensorFlow: v1.7.0, Horovod: 0.12.1, OpenMPI: 3.0.0. OS: CentOS 7.3, OpenMPU 23.0.0, Python 2.7.5
Time to Train to converge to 99% accuracy in model

intel

ARTIFICIAL INTELLIGENCE

SOLUTIONS
Solution Architects

AI Solutions Catalog (Public & Internal)

Platforms *Finance* *Healthcare* *Energy* *Industrial* *Transport* *Retail* *Home* *More...*

TOOLKITS
App Developers

DEEP LEARNING DEPLOYMENT

OpenVINO™
Open Visual Inference & Neural Network Optimization toolkit for inference deployment on CPU/GPU/FPGA for TF, Caffe & MXNet**

Intel® Movidius™ SDK
Optimized inference deployment on Intel VPUs for TensorFlow & Caffe**

REASONING

Intel® Saffron™ AI
Cognitive solutions on CPU for anti-money laundering, predictive maintenance, more

DEEP LEARNING

COMING SOON!

Intel® Deep Learning Studio
Open-source tool to compress deep learning development cycle

LIBRARIES
Data Scientists

MACHINE LEARNING LIBRARIES

Python
- Scikit-learn
- Pandas
- NumPy

R
- Cart
- Random Forest
- e1071

Distributed
- MlLib (on Spark)
- Mahout

DEEP LEARNING FRAMEWORKS

Now optimized for CPU

TensorFlow* MXNet* Caffe* BigDL/Spark*

COMING SOON!

Optimizations in progress

Caffe2* PyTorch* PaddlePaddle*

FOUNDATION
Library Developers

ANALYTICS, MACHINE & DEEP LEARNING PRIMITIVES

Python
Intel distribution optimized for machine learning

DAAL
Intel® Data Analytics Acceleration Library (incl machine learning)

MKL-DNN **clDNN**
Open-source deep neural network functions for CPU / integrated graphics

DEEP LEARNING GRAPH COMPILER

Intel® nGraph™ Compiler (Alpha)
Open-sourced compiler for deep learning model computations optimized for multiple devices (CPU, GPU, NNP) from multiple frameworks (TF, MXNet, ONNX)

HARDWARE
IT System Architects

AI FOUNDATION

Edge Device

DEEP LEARNING ACCELERATORS

COMING SOON!

Training

Inference

DATA-CENTRIC INNOVATION SUMMIT

ai.intel.com
#IntelDCISummit

WHERE DO WE GO FROM HERE?

HW-SW CO-EVOLUTION = IMPROVED AI DEVELOPER EXPERIENCE

NEURAL NETWORK

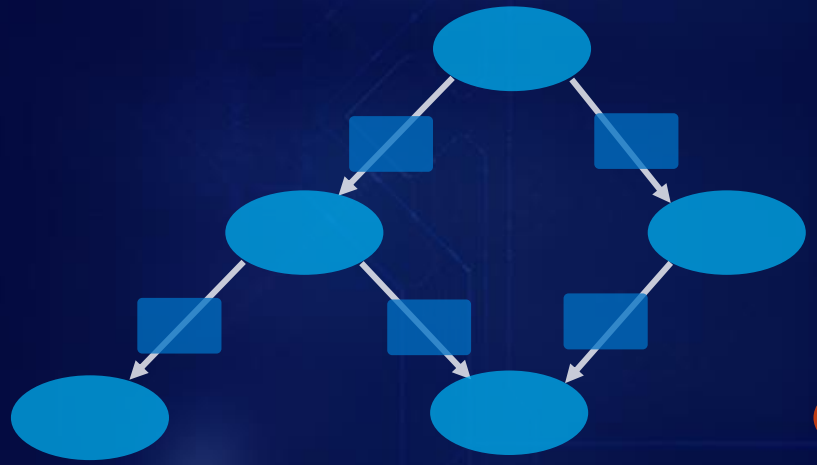

SUBGRAPHS

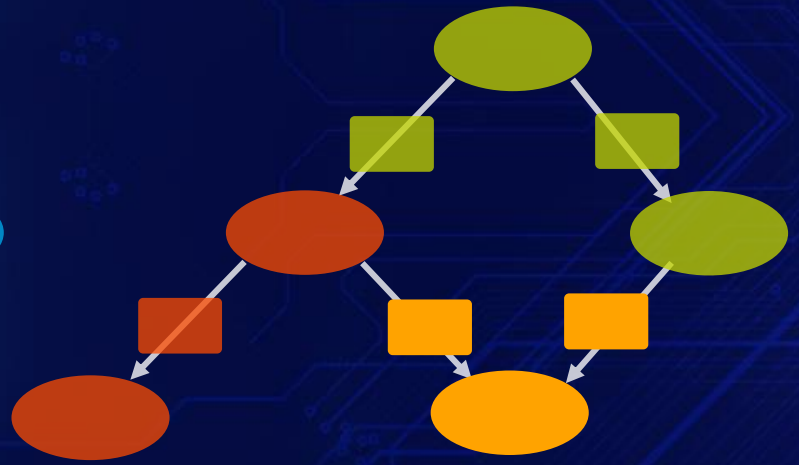

FED TO HARDWARE

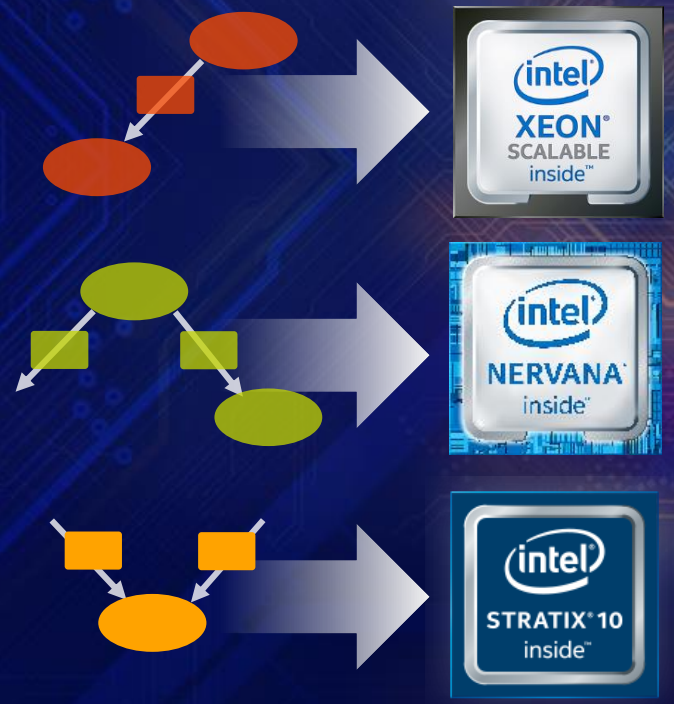

THE AI FUTURE

MORE MEMORY CLOSER TO COMPUTE = LARGER MODELS

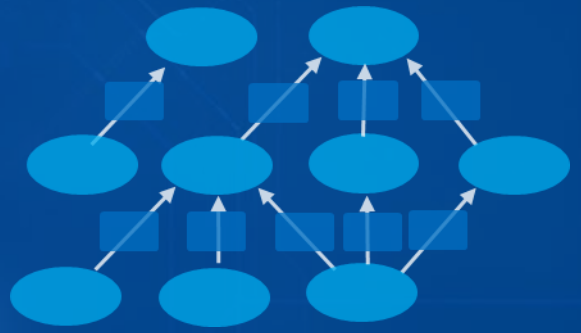

MIXED DATA TYPES = DENSER COMPUTE

BLURRING BETWEEN TRAINING AND INFERENCE = CONTINUOUS LEARNING

REQUIRES FLEXIBLE SOFTWARE AND HARDWARE DELIVERING WORKLOAD OPTIMIZED SOLUTIONS

Q&A

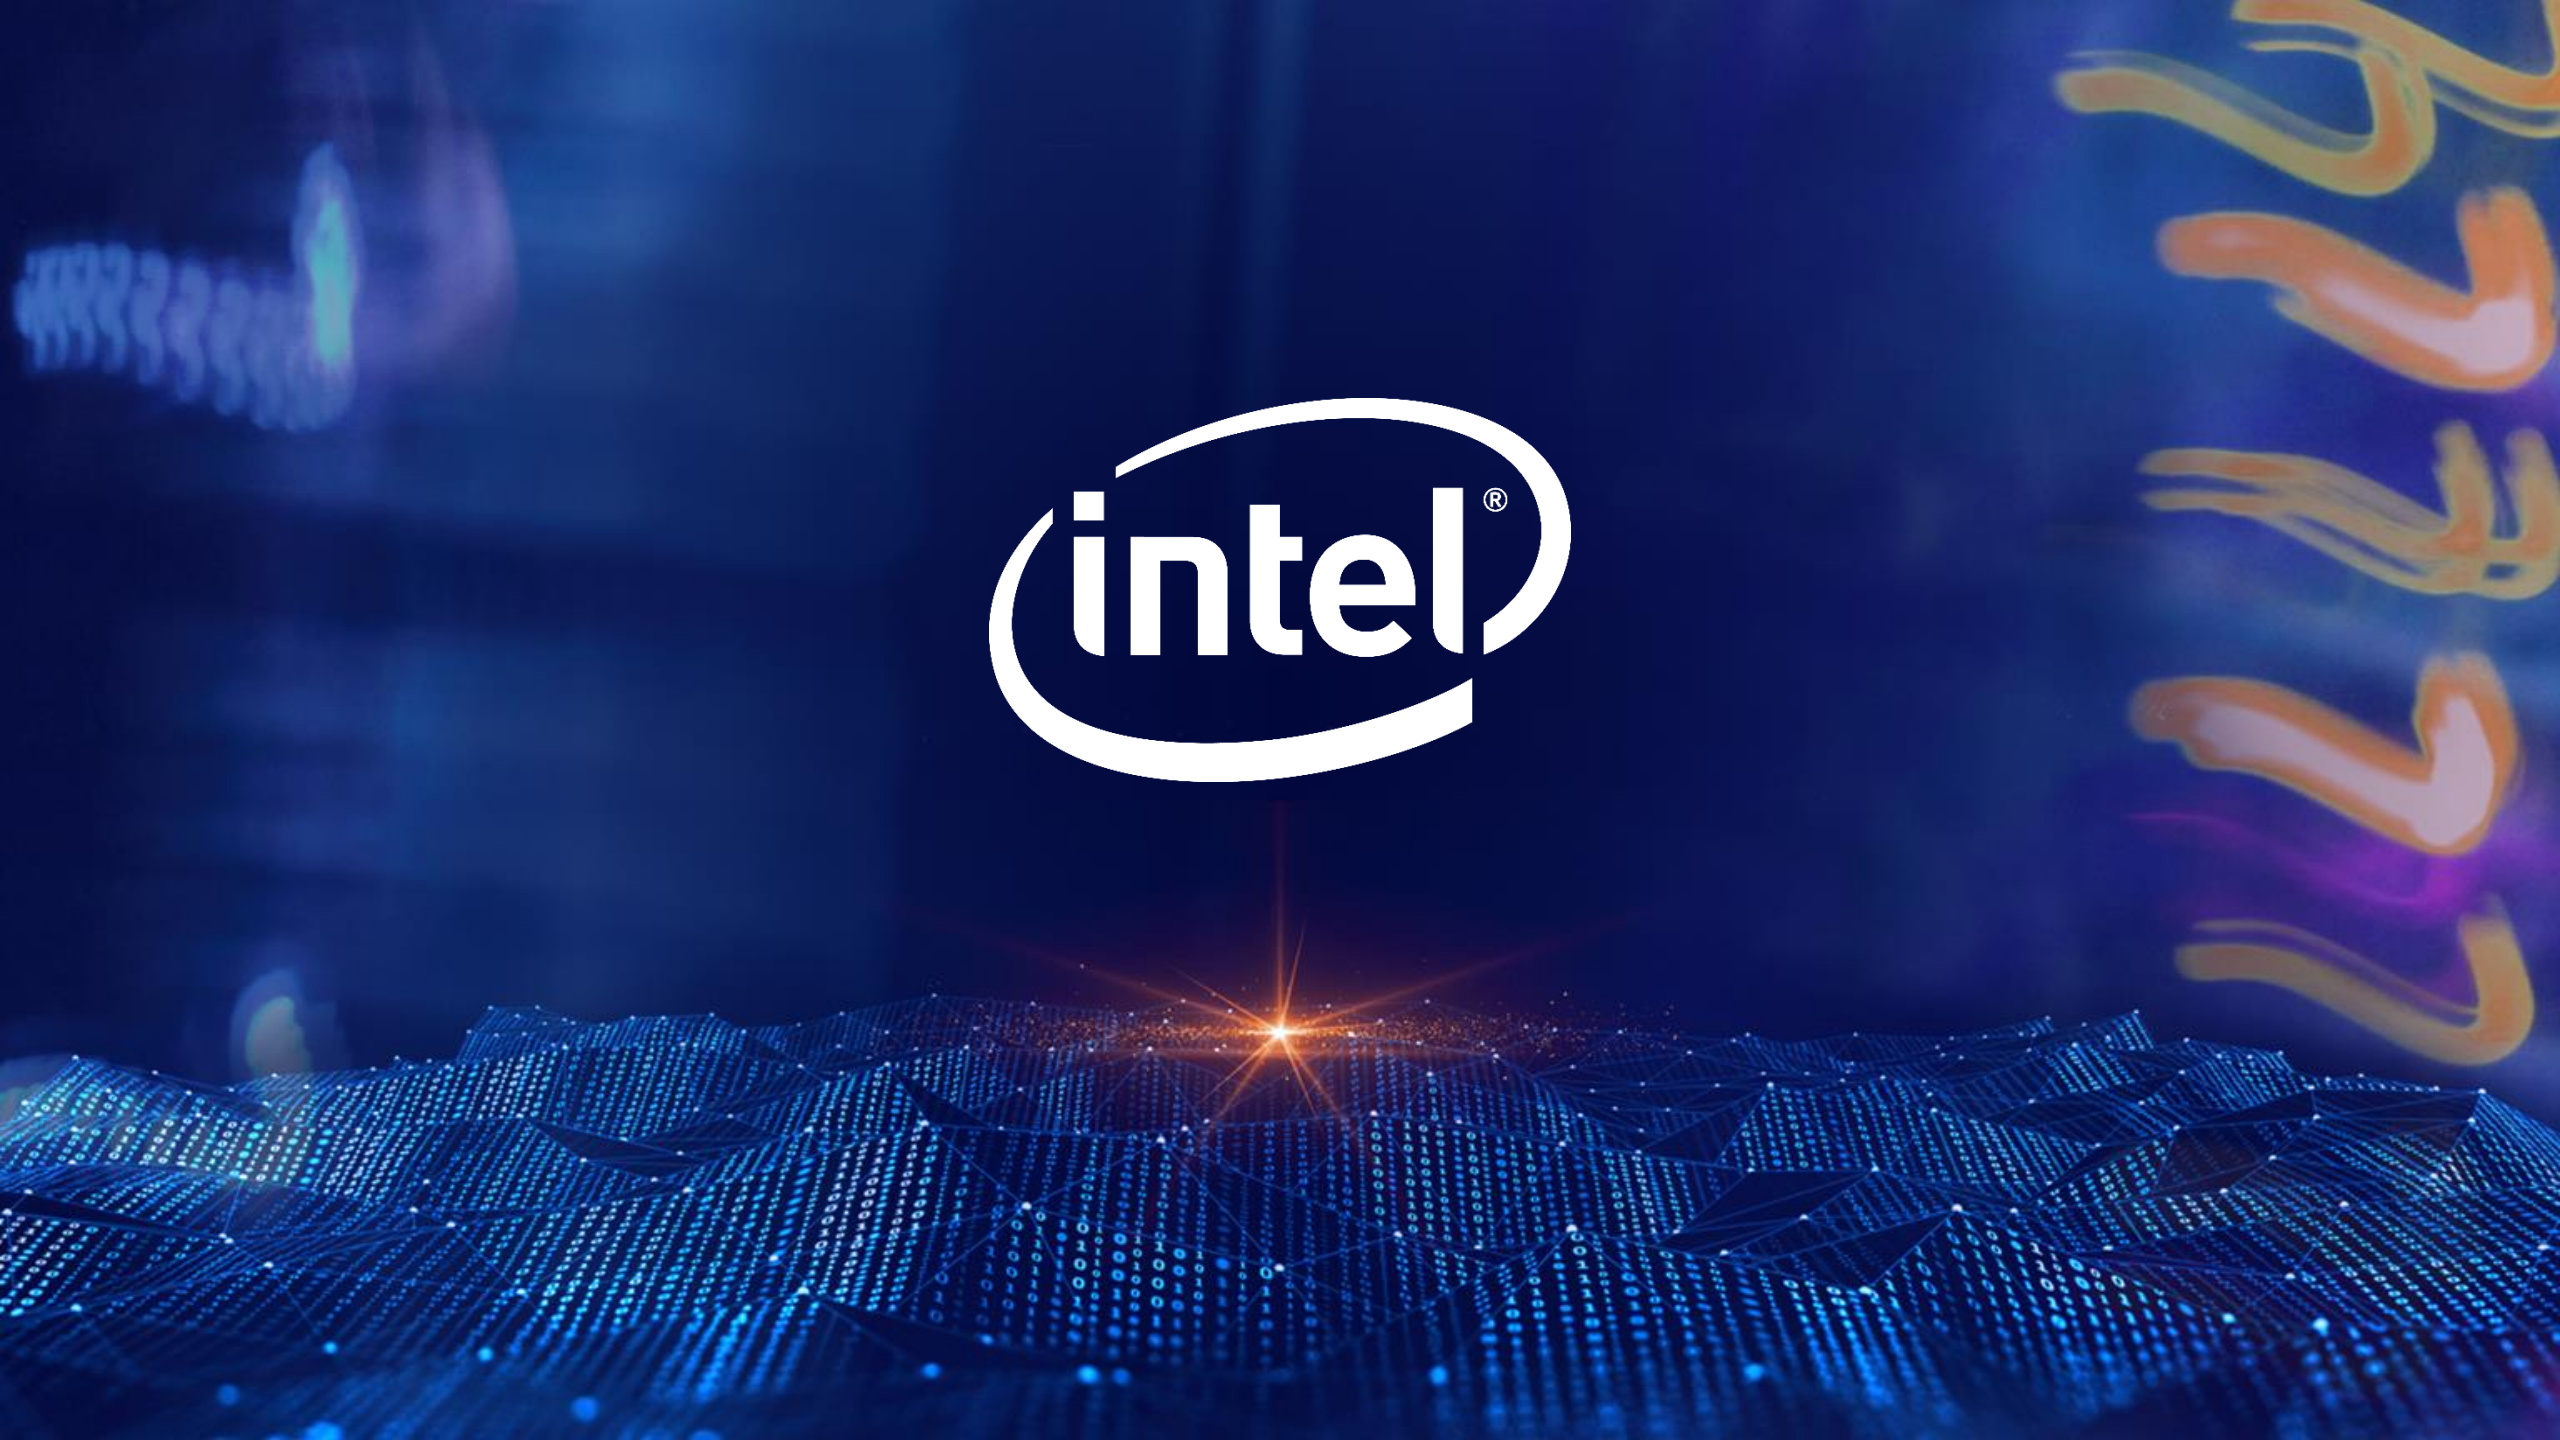
intel

NOTICES & DISCLAIMERS

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.**

Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. For more complete information about performance and benchmark results, visit http://www.intel.com/benchmarks .

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 http://www.intel.com/benchmarks .

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.

The benchmark results 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.

Intel does not control or audit third-party benchmark data or the web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.

DISCLOSURES

Statements in this presentation that refer to business outlook, future plans and expectations are forward-looking statements that involve a number of risks and uncertainties. Words such as "anticipates," "expects," "intends," "goals," "plans," "believes," "seeks," "estimates," "continues," "may," "will," “would,” "should," “could,” and variations of such words and similar expressions are intended to identify such forward-looking statements. Statements that refer to or are based on projections, uncertain events or assumptions also identify forward-looking statements. Such statements are based on management's current expectations, unless an earlier date is indicated, and involve many risks and uncertainties that could cause actual results to differ materially from those expressed or implied in these forward-looking statements. Important factors that could cause actual results to differ materially from the company's expectations are set forth in Intel's earnings release dated July 26, 2018, which is included as an exhibit to Intel’s Form 8-K furnished to the SEC on such date. Additional information regarding these and other factors that could affect Intel's results is included in Intel's SEC filings, including the company's most recent reports on Forms 10-K and 10-Q. Copies of Intel's Form 10-K, 10-Q and 8-K reports may be obtained by visiting our Investor Relations website at www.intc.com or the SEC's website at www.sec.gov.

All information in this presentation reflects management’s views as of the date of this presentation, unless an earlier date is indicated. Intel does not undertake, and expressly disclaims any duty, to update any statement made in this presentation, whether as a result of new information, new developments or otherwise, except to the extent that disclosure may be required by law.