
DATA-CENTRIC
INNOVATION DAY
MOVE | STORE | PROCESS


intel
NAVIN SHENOY
EXECUTIVE VICE PRESIDENT & GENERAL MANAGER
DATA CENTER GROUP

THE DATA-CENTRIC WORLD

OVER **HALF** OF THE WORLD'S DATA

WAS CREATED IN THE LAST **2 YEARS**

LESS THAN **2%** HAS BEEN ANALYZED

#datacentric | intel

Source: Data Age 2025, sponsored by Seagate with data from IDC Global DataSphere, Nov 2018

INDUSTRY MEGA TRENDS

PROLIFERATION OF
CLOUD COMPUTING

GROWTH OF
AI & ANALYTICS

CLOUDIFICATION OF THE
NETWORK & EDGE

#datacentric | intel

EXPLOSION IN DEMAND FOR COMPUTE

INCREASING COMPUTE DEMAND
DIVERSIFYING WORKLOAD NEEDS

- AI
- ANALYTICS
- HPC
- MULTI-CLOUD & ORCHESTRATION
- NETWORK
- DATABASE
- VIRTUALIZATION
- SECURITY

Source: Intel analysis

#datacentric | intel

DATA-CENTRIC INFRASTRUCTURE FOCUS

MOVE FASTER

- intel ETHERNET
- intel SILICON PHOTONICS
- intel OMNI-PATH FABRIC

STORE MORE

- intel OPTANE™ DC PERSISTENT MEMORY
- intel OPTANE™ DC SOLID STATE DRIVE

PROCESS EVERYTHING

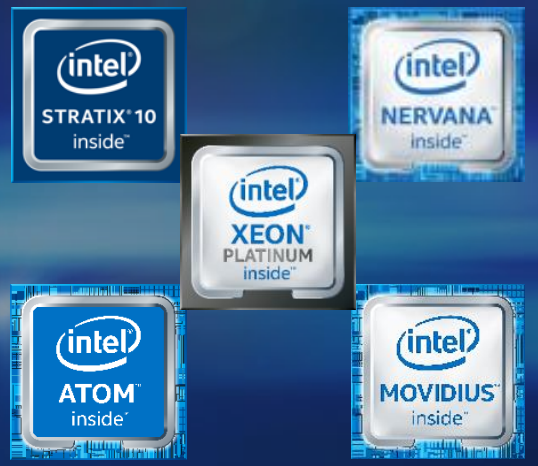

SOFTWARE & SYSTEM-LEVEL OPTIMIZED

#datacentric | intel

LAUNCHING TODAY

DATA-CENTRIC PORTFOLIO

MOVE FASTER

INTEL® ETHERNET 800 SERIES ADAPTER

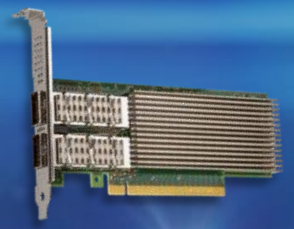

STORE MORE

INTEL® SSDs

INTEL® OPTANE™ DC PERSISTENT MEMORY

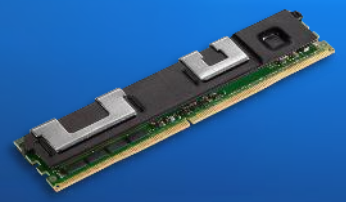

PROCESS EVERYTHING

2ND GENERATION INTEL® XEON® SCALABLE

INTEL® XEON® D-1600

INTEL® AGILEX™

intel select solution

#datacentric | intel


OUR PARTNERS
1&1 IONOS
AAEON
accenture
ADLINK Leading EDGE COMPUTING
ADVANTECH
AEROSPIKE
AEWIN
affirmed
AGENT
AIC
Alibaba Group 阿里巴巴集团
ALTIBASE
AMAX
Apache CASSANDRA
Apache Spark
ASA
AsiaInfo 亚信科技
at&t
atrida
Atos
avanade
aws
AXIOMTEK
Baidu Cloud
宝信软件 BAOSIGHT
bluechip
BONC 东方国信
BOSTON Servers | Storage | Solutions
ByteDance 字节跳动
CASWELL Complete Your Solutions
CDW
Celestica
Check Point SOFTWARE TECHNOLOGIES LTD
CIARA
cisco
cloudera
云从科技 CLOUDWALK
COMCAST
COLFAX Customized Solutions
CRAY
criteo
alhua TECHNOLOGY
databricks
dataON DataONSTORAGE
DATASTAX
DELL EMC
DFI
EISOO
LOM MICROSYSTEMS 珑微系统
安擎 NGINETECH
EQUUS COMPUTE SOLUTIONS
ERICSSON
f5
FiberHome
FORTINET
FUJITSU
GBASE
GIGABYTE
GIGASPACES innovate with confidence
Google Cloud
H2O.ai
H3C
hazelcast
Hewlett Packard Enterprise
海鑫科金 HISIGN TECHNOLOGY
HITACHI
HLRN
HUAWEI
HYTRUST
IBM
iEi
科大讯飞 iFLYTEK
Insight
inspur
Inventec
JD.COM
JABIL
JUNIPER NETWORKS
人大金仓 Kingbase
Kingsoft Cloud KSCLOUD
kontron
kx
Lanner
Lenovo
LOCKHEED MARTIN
LOGIC SUPPLY
宏杉科技 macrosan
MEGWARE
mercury systems
Microsoft
南京云创大数据科技股份有限公司 Nanjing Innovative Data Technologies,Inc.
NARI 国电南瑞 NARI-TECH
ncs
NetApp
Neusoft
NEXCOM
NGINX
NOKIA
NOR-TECH
NUTANIX
online.net
ORACLE
OVH.com
PENGUIN COMPUTING
phoenixNAP GLOBAL IT SERVICES
宝德 PowerLeader
PRESIDIO
QCT
QNAP
Qwilt
Rakuten
redhat
redislabs
RENIAC
RocksDB
RSC
SAP
sas
SEAGATE
天玑科技 Pioneer in IT Service
shi
SIEMENS Healthineers
Silicom Ltd. Connectivity Solutions
中科曙光 Sugon
SUNJESOFT SUNJESOFT Inc.
Super Cluster Inc. 超集信息科技有限公司
SUPERMICRO
SUSE
Tencent Cloud
THINKMATE HIGH PERFORMANCE COMPUTING
TRANSWARP
TRENTON SYSTEMS
TYAN
ubuntu
UCLOUD
UNICOM Engineering, Inc. A Division of UNICOM Global
Virtuozzo
vmware
vodafone
wiwynn
World Wide Technology
WORTMANN AG
Yandex
ZTE
#datacentric
intel

2ND GENERATION
INTEL® XEON® SCALABLE PROCESSORS

>50	DOZENS	8 TO 56	4.5TB	1 TO 8
STANDARD SKUS	CUSTOM SKUS	CORES PER SOCKET	MEMORY PER SOCKET	SOCKETS

- INTEL® OPTANE™ DC PERSISTENT MEMORY
- INTEL® DEEP LEARNING BOOST
- INTEL® SPEED SELECT TECHNOLOGY
- NETWORK-OPTIMIZED SKUS
- CLOUD-OPTIMIZED SKUS
- SECURITY MITIGATIONS

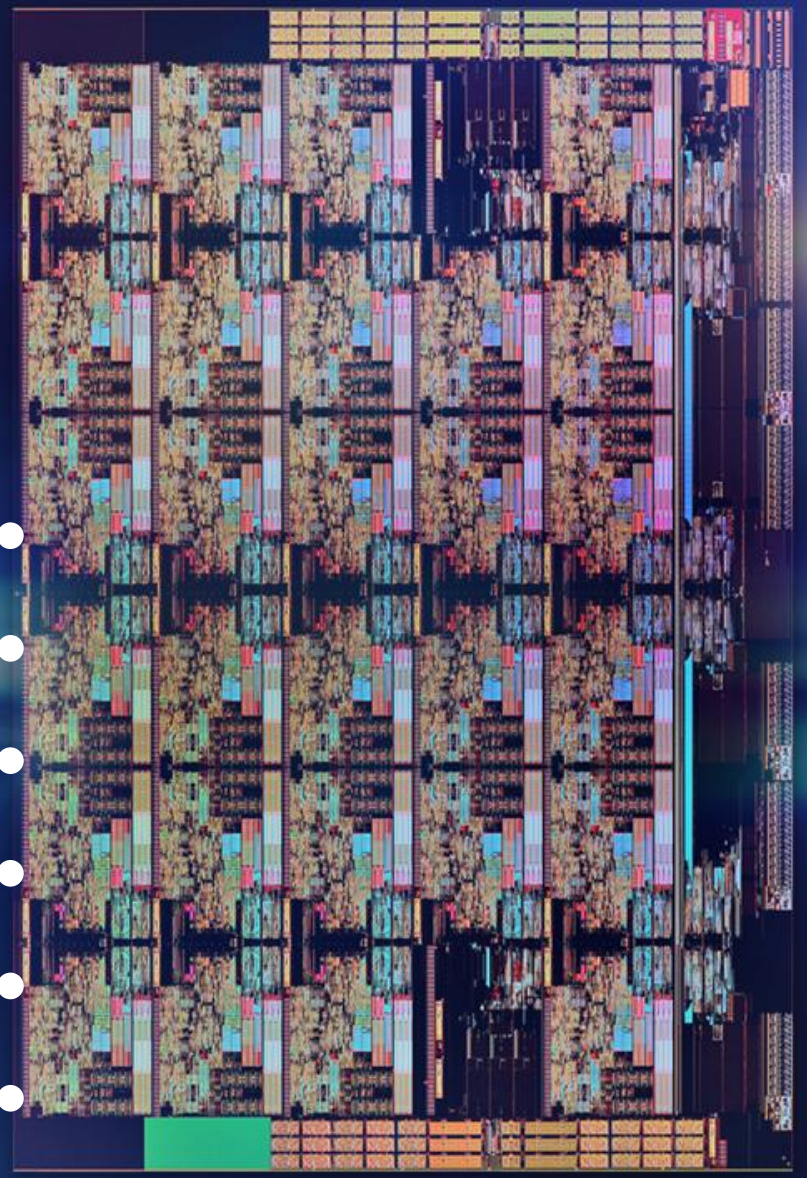

BUILDING ON **20 YEARS** OF DATA CENTER PROCESSOR INNOVATION

#datacentric | intel

ALL NEW LEVEL OF ADVANCED PERFORMANCE

INTEL® XEON® PLATINUM 9200 PROCESSORS

HIGHEST PERFORMANCE

UP TO 56 CORES
PER SOCKET

LEADERSHIP

PERFORMANCE PER RACK

HIGHEST BANDWIDTH

UP TO 12 CHANNELS
NATIVE DDR4 MEMORY

DESIGNED FOR THE MOST **DATA-INTENSIVE WORKLOADS**

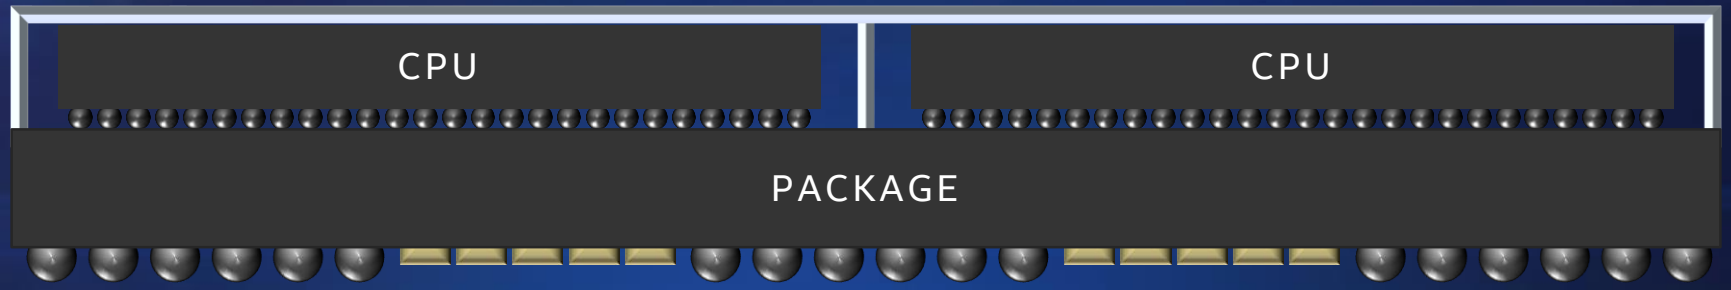

#datacentric | intel

Leadership performance per rack based on 4 benchmarks (Integer Throughput, Floating Point Throughput, Memory Bandwidth and LINPACK).


HLRN
#datacentric
intel

WORLD RECORD + REAL WORKLOAD
PERFORMANCE LEADERSHIP

LAMMPS
BUSINESS ANALYTICS
2.19X
9242 VS 8160

LS-DYNA
BUSINESS ANALYTICS
2.01X
9242 VS 8160

宝信软件 BAOSIGHT
BUSINESS ANALYTICS
1.39X
8280+OPTANE PM VS DRAM

AsiaInfo 亚信科技
BUSINESS ANALYTICS
1.54X
8280+OPTANE PM VS DRAM

MAXIMIZING MAINSTREAM SKUS

UP TO **1.33X** AVERAGE PERF GAIN GEN ON GEN

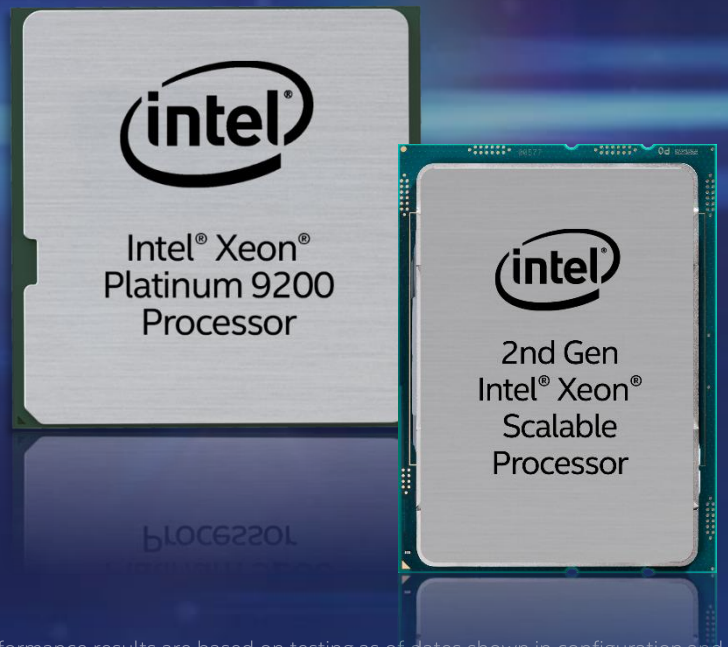

HUAWEI
CLOUD MANAGEMENT
1.42X MORE VMS
8260+OPTANE PM VS DRAM

GBASE®
IN-MEMORY DATABASE
1.35X
8260+OPTANE PM VS DRAM

云从科技 CLOUDWALK
BUSINESS ANALYTICS
3.38X LOWER LATENCY
8260 DLBOOST VS FP32

海鑫科金 HISIGN TECHNOLOGY
BUSINESS ANALYTICS
2.19X
8260 DLBOOST VS FP32

NOKIA
vNETWORK GATEWAY
2.0X
5218N+QAT VS 5118

#datacentric intel

AI: THE FASTEST GROWING WORKLOAD

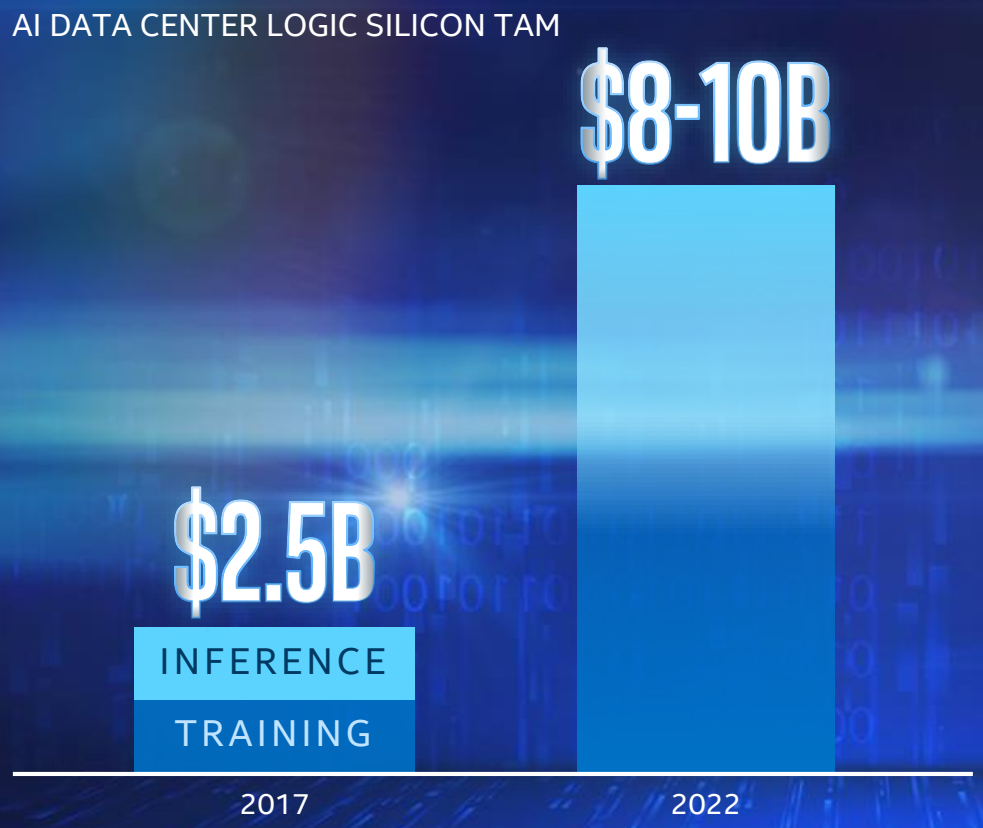

INFERENCE >50% OF THE AI SI OPPORTUNITY

Source: AI Si 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.

INTEL® DEEP LEARNING BOOST

OPTIMIZING AI INFERENCE

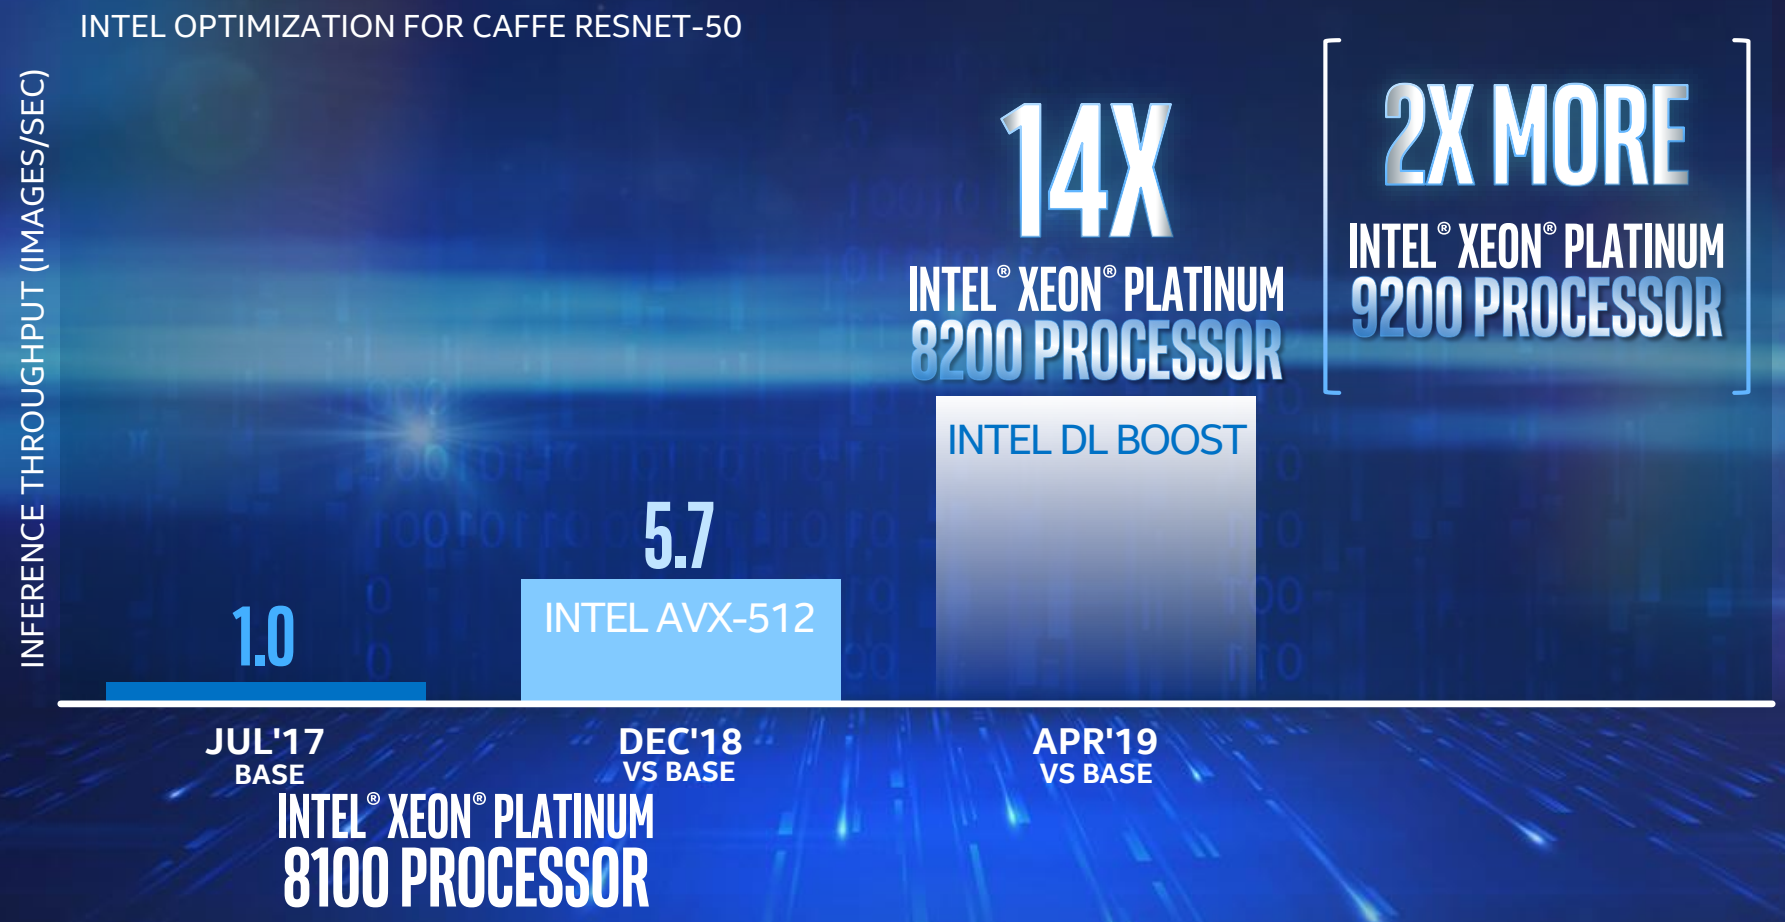

#datacentric | intel

MAKING AI ON IA EASY

INTEL® DL BOOST ECOSYSTEM SUPPORT

Microsoft
3.4X
IMAGE RECOGNITION

JD.COM
2.4X
TEXT DETECTION

Alibaba Group 阿里巴巴集团
2-4X
8 DIFFERENT WORKLOADS

4.43X
ML INFERENCING

Tencent Cloud
3.26X
VIDEO ANALYSIS

OPTIMIZED SW & FRAMEWORKS	Caffe, mxnet, OpenVINO, ONNX, PaddlePaddle, PyTorch, TensorFlow
SOFTWARE VENDORS	AsiaInfo, 云从科技 CLOUDWALK, H2O.ai, 海鑫科金 HISIGN TECHNOLOGY, Neusoft
CLOUD SERVICE PROVIDERS	Alibaba Group 阿里巴巴集团, aws, Baidu Cloud, JD.COM, Microsoft, Tencent Cloud
ENTERPRISES	alhua TECHNOLOGY, 科大讯飞 iFLYTEK, SIEMENS Healthineers


aws
MATT GARMAN
VICE PRESIDENT
AWS COMPUTE SERVICES
#datacentric
intel

CLOUDIFICATION OF THE NETWORK

DATA CENTER | CLOUD | CORE | ACCESS | EDGE | DEVICES | THINGS

THE NETWORK MOVES TO IA

2011 — **NFV** DEFINED

2013 — **1ST NFV** PROOF OF CONCEPTS

2015 — **20%** OF COMMS SPS ADOPT NFV

2017 — **DPDK** MOVES TO LINUX FOUNDATION

2018 — **65%** NETWORK IS VIRTUALIZED — at&t

2019 — **1ST 100%** CLOUD-NATIVE NETWORK — Rakuten

#datacentric | intel

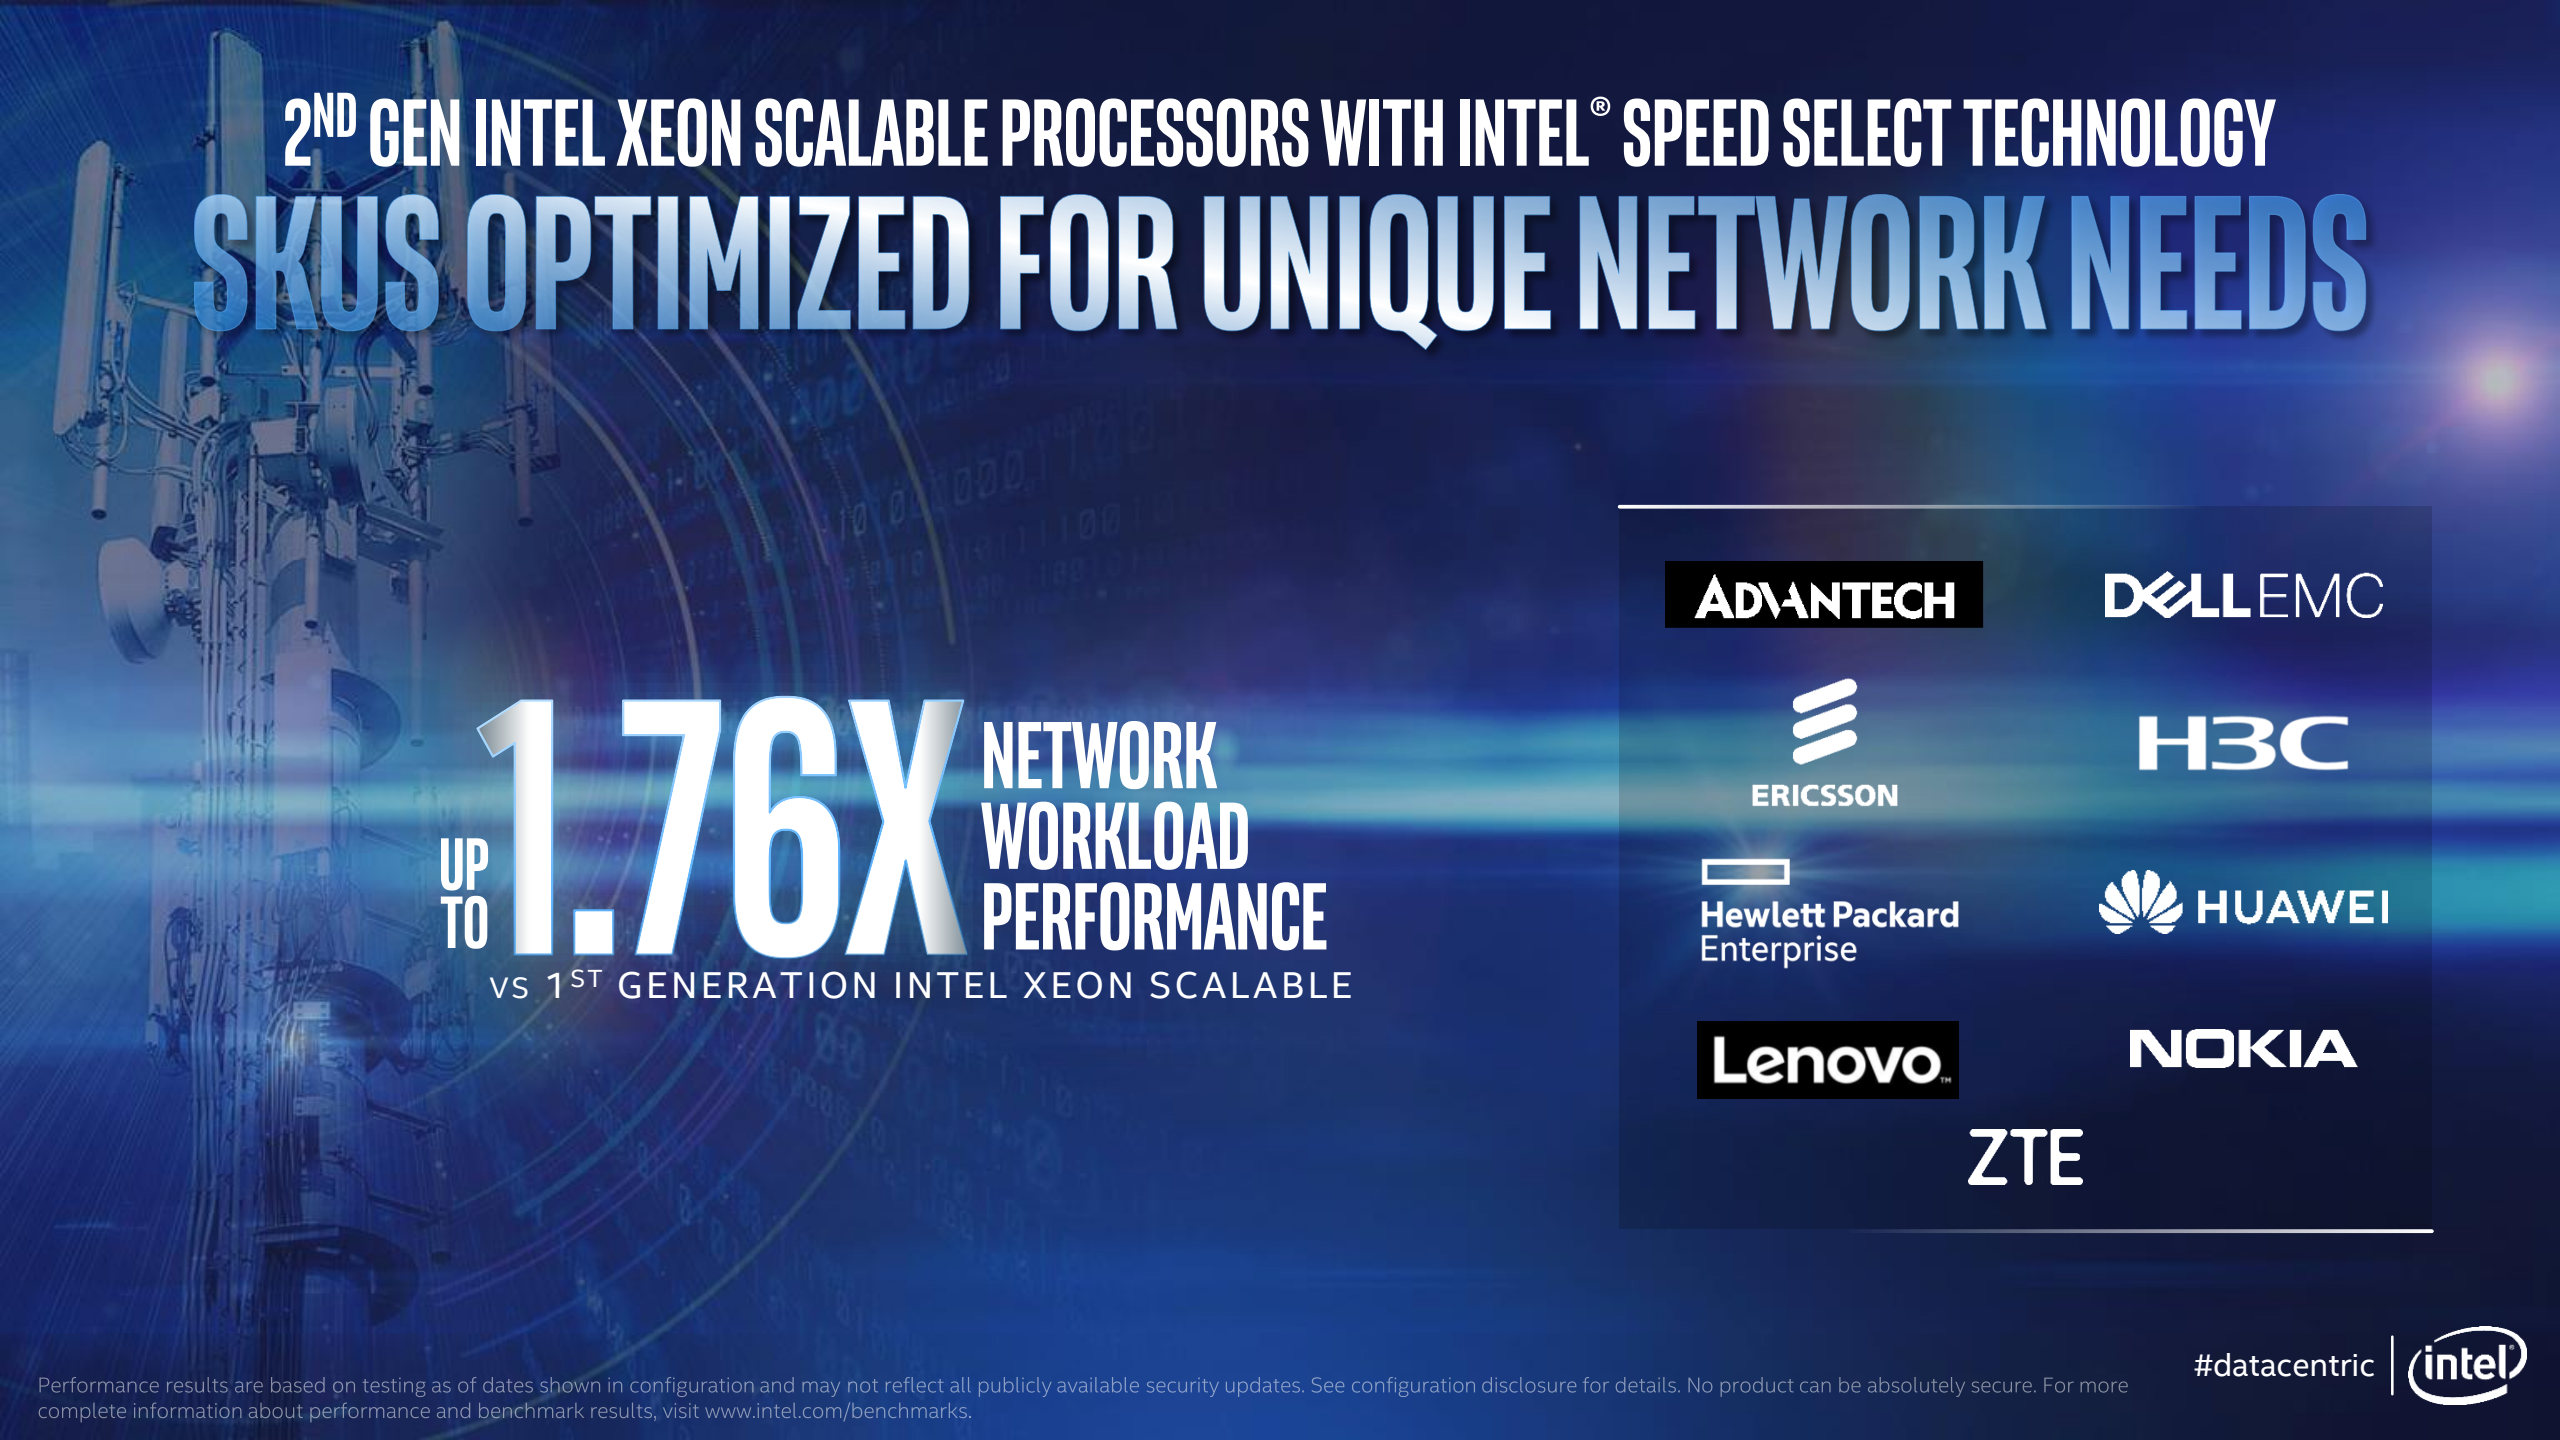

Performance results are based on testing as of dates shown in configuration and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks.


vodafone
#datacentric
intel


SYSTEM-LEVEL
OPTIMIZATION
UNLEASHES
PERFORMANCE
#datacentric
intel

@mattbytes

MATT SINGER

SR. STAFF HARDWARE ENGINEER
TWITTER

#datacentric | intel

LAUNCHING TODAY

DATA-CENTRIC PORTFOLIO

MOVE FASTER

INTEL® ETHERNET 800 SERIES ADAPTER

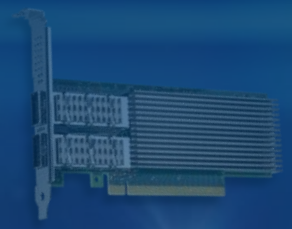

STORE MORE

INTEL® SSDs

INTEL® OPTANE™ DC PERSISTENT MEMORY

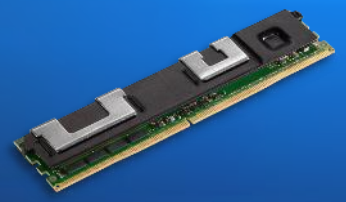

PROCESS EVERYTHING

2ND GENERATION INTEL® XEON® SCALABLE

INTEL® XEON® D-1600

INTEL® AGILEX™

intel select solution

#datacentric | intel

REMOVING STORAGE BOTTLENECKS

DUAL-PORT
INTEL® OPTANE™ DC SSD D4800X

24X7 ACCESS

FOR MISSION CRITICAL ENTERPRISE STORAGE

INTEL® SSD D5-P4326
QLC NAND 'RULER'

UP TO 20X

STORAGE RACK CONSOLIDATION
VS HARD DRIVES

#datacentric | intel

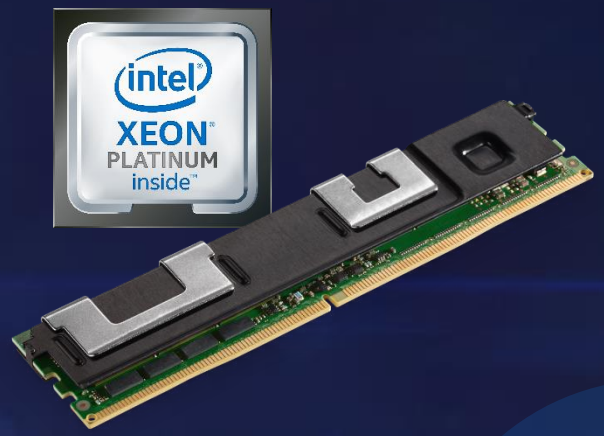

INTEL® OPTANE™ DC PERSISTENT MEMORY

MEMORY INNOVATION 10 YEARS IN THE MAKING

ECOSYSTEM SUPPORT

SOLUTION OPTIMIZATION

TECHNOLOGY INNOVATIONS

UP TO 36TB
8 SOCKET SYSTEM

9.1B SAP BW ON HANA RECORDS
NEW WORLD RECORD

redis
UP TO 8X MORE VM INSTANCES
MEETING SUB-mS SLA

#datacentric intel


SAP
DIRK BASENACH
SENIOR VICE PRESIDENT
HEAD OF DATABASE
#datacentric
intel

intel OPTANE™ DC PERSISTENT MEMORY ECOSYSTEM SUPPORT

#datacentric | intel


Google Cloud
BART SANO
VICE PRESIDENT
PLATFORMS
#datacentric
intel

LAUNCHING TODAY

DATA-CENTRIC PORTFOLIO

MOVE FASTER

INTEL® ETHERNET 800 SERIES ADAPTER

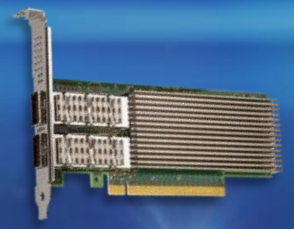

STORE MORE

INTEL® SSDs

INTEL® OPTANE™ DC PERSISTENT MEMORY

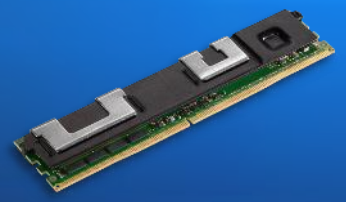

PROCESS EVERYTHING

2ND GENERATION INTEL® XEON® SCALABLE

INTEL® XEON® D-1600

INTEL® AGILEX™

intel select solution

#datacentric | intel


intel

SYSTEM CONFIGURATION: LEADERSHIP PERFORMANCE PER RACK

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.

Performance per rack leadership based on 4 benchmarks (Integer Throughput, Floating Point Throughput, Memory Bandwidth and LINPACK). Details below:
Integer Throughput:
1-node, 2x Intel® Xeon® Platinum 9282 processor on Walker Pass with 768 GB (24x 32GB 2933) total memory, ucode 0x4000010 on CentOS Linux release 7.6.1810, 4.20.0+, IC19u1, AVX512, HT on, Turbo on, score: est int throughput=628, test by Intel on 3/14/2019. Rack performance estimate of 40192. 42U rack, 32U dedicated to compute, total of 64 compute nodes. 64 * 628 = 40192
1-node, 2x AMD* EPYC* 7601, https://www.spec.org/cpu2017/results/res2019q1/cpu2017-20190304-11124.html, score: 301, test by Dell on Feb 2019 Rack performance estimate of 19264. 42U rack, 32U dedicated to compute, total of 64 compute nodes. 64 * 301 = 19264

Floating Point Throughput:
1-node, 2x Intel® Xeon® Platinum 9282 processor on Walker Pass with 768 GB (24x 32GB 2933) total memory, ucode 0x4000010 on CentOS Linux release 7.6.1810, 4.20.0+, IC19u1, AVX512, HT on, Turbo on, score: est fp throughput=522, test by Intel on 3/14/2019. Rack performance estimate of 33408. 42U rack, 32U dedicated to compute, total of 64 compute nodes. 64 * 522 = 33408
1-node, 2x AMD* EPYC* 7601, https://www.spec.org/cpu2017/results/res2019q1/cpu2017-20190304-11125.html, score: 282, test by Dell on Feb 2019 Rack performance estimate of 17152. 42U rack, 32U dedicated to compute, total of 64 compute nodes. 64 * 282 = 18048

Memory Bandwidth:
1-node, 2x Intel® Xeon® Platinum 9282 processor on Walker Pass with 768 GB (24x 32GB 2933) total memory, ucode 0x4000010, on CentOS Linux release 7.6.1810, 4.20.0+, IC19u1, AVX512, HT off, Turbo on, score: Stream Triad=407GiB/s, test by Intel on 3/14/2019. Rack performance estimate of 26048. 42U rack, 32U dedicated to compute, total of 64 compute nodes. 64 * 407 = 26048
1-node, 2x AMD* EPYC* 7601, http://www.amd.com/system/files/2017-06/AMD-EPYC-SoC-Delivers-Exceptional-Results.pdf, score=290, test by AMD as of June 2017. Rack performance estimate of 18560. 42U rack, 32U dedicated to compute, total of 64 compute nodes. 64 * 290 = 18560

LINPACK:
1-node, 2x Intel® Xeon® Platinum 9282 processor on Walker Pass with 768 GB (24x 32GB 2933) total memory, ucode 0x4000010 on CentOS Linux release 7.6.1810, 4.20.0+, IC19u1, N 210000, AVX512 MKL 2019, HT off, Turbo on, score: Intel® Distribution of LINPACK=6411, test by Intel on 3/14/2019. Rack performance estimate of 410.3 TFlops. 42U rack, 32U dedicated to compute, total of 64 compute nodes. 64 * 6411 = 410.3 TFlops
1-node, 2x AMD EPYC 7601: Supermicro AS-2023US-TR4 with 2 AMD EPYC 7601 (2.2GHz, 32 core) processors, SMT OFF, Turbo ON, BIOS ver 1.1a, 4/26/2018, microcode: 0x8001227, 16x32GB DDR4-2666, 1 SSD, Ubuntu 18.04.1 LTS (4.17.0-041700-generic Retpoline), High Performance Linpack v2.2, compiled with Intel(R) Parallel Studio XE 2018 for Linux, Intel MPI version 18.0.0.128, AMD BLIS ver 0.4.0, Benchmark Config: Nb=232, N=168960, P=4, Q=4, Score =1095GFs, tested by Intel as of July 31, 2018. Rack performance estimate of 70.08 TFlops. 42U rack, 32U dedicated to compute, total of 64 compute nodes. 64 * 1095 = 70.08 TFLOPs

SYSTEM CONFIGURATION: WORLD RECORD + REAL WORKLOAD PERFORMANCE LEADERSHIP

Performance results are based on testing as of dates shown in configuration and may not reflect all publicly available security updates. See configuration disclosure for details. No product or component 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 visit www.intel.com/benchmarks.

2.19x LAMMPS* Water: 1-node, 2x Intel® Xeon® Platinum 8160L cpu on Wolf Pass with 192 GB (12 slots / 16GB / 2666) total memory, ucode 0x200004d on Oracle Linux Server release 7.6 , 3.10.0-862.14.4.el7.crt1.x86_64, Intel SSDSC2BA80, LS-Dyna 9.3-Explicit AVX2 binary, 3car, HT on, Turbo on, test by Intel on 2/26/2019. 1-node, 2x Intel® Xeon® Platinum 9242 cpu on Intel reference platform with 384 GB (24 slots / 16GB / 2933) total memory, ucode 0x4000017 on CentOS 7.6, 3.10.0-957.5.1.el7.x86_64, Intel SSDSC2BA80, LS-Dyna 9.3-Explicit AVX2 binary, 3car, HT on, Turbo on, test by Intel on 3/18/2019.

2.01x LS-Dyna* Explicit, 3car: 1-node, 2x Intel® Xeon® Platinum 8160L cpu on Wolf Pass with 192 GB (12 slots / 16GB / 2666) total memory, ucode 0x200004d on Oracle Linux Server release 7.6 , 3.10.0-862.14.4.el7.crt1.x86_64, Intel SSDSC2BA80, LAMMPS version 12 Dec 2018, Water, HT on, Turbo on, test by Intel on 2/26/2019. 1-node, 2x Intel® Xeon® Platinum 9242 cpu on Intel reference platform with 384 GB (24 slots / 16GB / 2933) total memory, ucode 0x4000017 on CentOS 7.6, 3.10.0-957.5.1.el7.x86_64, Intel SSDSC2BA80, LAMMPS version 12 Dec 2018, Water, HT on, Turbo on, test by Intel on 3/8/2019.

1.39x BAOSIGHT* xInsight*: 1-node, 2x Intel® Xeon® Platinum 8260L cpu on S2600WFS with 768 DDR GB (24 slots / 32GB / 2666) total memory, ucode 0x400000A on CentOS 7.5, 3.10.0-957.1.3.el7.x86_64, 1x Intel 480GB SSD OS Drive, 1x Intel XC722, xInsight 2.0 internal workload, HT on, Turbo on, test by Intel/Baosight on 1/8/2019. 1-node, 2x Intel® Xeon® Platinum 8260L cpu on S2600WFS with 192 DDR + 1024 Intel DCPMM GB (12 slots / 16 GB / 2666 DDR + 8 slots / 128 GB / 2666 Intel DCPMM) total memory, ucode 0x400000A on CentOS 7.5, 3.10.0-957.1.3.el7.x86_64, 1x Intel 480GB SSD OS Drive, 1x Intel XC722, xInsight 2.0 internal workload, HT on, Turbo on, test by Intel/Baosight on 1/9/2018.

1.54x AsiaInfo* BSS*: 1-node, 2x Intel® Xeon® Platinum 8180 cpu on S2600WFD with 768 GB (24 slots / 32GB / 2666) total memory, ucode 0x2000035 on RedHat 7.5, 3.10.0-957.1.3.el7.x86_64, 1x Intel 400GB SSD OS Drive, 1 x P4500 1TB Application Data, 1x Intel XC722, BSS 3.1.1 + self defined workload, HT on, Turbo on, test by Intel/AsiaInfo on 12/27/2018. 1-node, 2x Intel® Xeon® Platinum 8280 cpu on S2600WFD with 192 DDR + 1024 Intel DCPMM GB (12 slots / 16 GB / 2666 DDR + 8 slots / 128 GB / 2666 Intel DCPMM) total memory, ucode 0x400000A on RedHat 7.5, 3.10.0-957.1.3.el7.x86_64, 1x Intel 400GB SSD OS Drive, , 1x Intel XC722, BSS 3.1.1 + self defined workload, HT on, Turbo on, test by Intel/AsiaInfo on 12/26/2018.

1.42x Huawei* FusionSphere*: 1-node, 2x Intel® Xeon® Platinum 8260L cpu on Wolf Pass with 1024 GB (16 slots / 64GB / 2666) total memory, ucode 0x400000A on FusionSphere HyperV, 3.10.0-514.44.5.10_96.x86_64 , 1x Intel 800GB SSD OS Drive, 1x Intel 800GB SSD OS Drive, 1x Intel XC722, FusionSphere 6.3.1, mysql-5.7.24, sysbench-1.0.6, HT on, Turbo on, test by Huawei/Intel on 1/11/2018. 1-node, 2x Intel® Xeon® Platinum 8260L cpu on Wolf Pass with 384 DDR + 1536 Intel DCPMM GB (12 slots / 32 GB / 2666 DDR + 12 slots / 128 GB / 2666 Intel DCPMM) total memory, ucode 0x400000A on FusionSphere HyperV, 3.10.0-514.44.5.10_96.x86_64 , 3 x P3520 1.8TB Application Data, 3 x P3520 1.8TB Application Data, 1x Intel XC722, FusionSphere 6.3.1, mysql-5.7.24, sysbench-1.0.6, HT on, Turbo on, test by Huawei/Intel on 1/11/2018.

1.35x GBASE: 1-node, 2x Intel® Xeon® Platinum 8260 cpu on S2600WFT with 768 DDR GB (24 slots / 32GB / 2666) total memory, ucode 0x400000A on CentOS 7.5, 3.10.0-957.1.3.el7.x86_64, 1x Intel 400GB SSD OS Drive, 1x Intel XC722, Gbase 8m 6.3.2 OCS Benchmark, HT on, Turbo on, test by GBASE/Intel on 2/19/2019. 1-node, 2x Intel® Xeon® Platinum 8260 cpu on S2600WFT with 192 DDR + 1024 Intel DCPMM GB (12 slots / 16 GB / 2666 DDR + 8 slots / 128 GB / 2666 Intel DCPMM) total memory, ucode 0x400000A on CentOS 7.5, 3.10.0-957.1.3.el7.x86_64, 1x Intel 400GB SSD OS Drive, 1x Intel XC722, Gbase 8m 6.3.2 OCS Benchmark, HT on, Turbo on, test by GBASE/Intel on 2/19/2019.

Up to 1.33x average generational gains on mainstream Gold SKU: Geomean of est SPECrate2017_int_base, est SPECrate2017_fp_base, Stream Triad, Intel Distribution of Linpack, server side Java. Gold 5218 vs Gold 5118: 1-node, 2x Intel® Xeon® Gold 5218 cpu on Wolf Pass with 384 GB (12 X 32GB 2933 (2666)) total memory, ucode 0x4000013 on RHEL7.6, 3.10.0-957.el7.x86_64, IC18u2, AVX2, HT on all (off Stream, Linpack), Turbo on, result: est int throughput=162, est fp throughput=172, Stream Triad=185, Linpack=1088, server side java=98333, test by Intel on 12/7/2018. 1-node, 2x Intel® Xeon® Gold 5118 cpu on Wolf Pass with 384 GB (12 X 32GB 2666 (2400)) total memory, ucode 0x200004D on RHEL7.6, 3.10.0-957.el7.x86_64, IC18u2, AVX2, HT on all (off Stream, Linpack), Turbo on, result: est int throughput=119, est fp throughput=134, Stream Triad=148.6, Linpack=822, server side java=67434, test by Intel on 11/12/2018.

3.38x Cloudwalk inference latency improvement: 1-node, 2x Intel Xeon Platinum 8260L cpu on S2600WFS with 192 GB (12 slots / 16 GB / 2666 MHz) total memory, ucode 0x400000A on CentOS 7.5, 3.10.0-957.1.3.el7.x86_64, 1x Intel 480GB SSD OS Drive, 1 x P4500 1TB Application Data, 1x Intel XC722, Cloudwalk Facial Recognition, GCC 4.8.5, Intel MKL-DNN, Intel Optimization for Caffe 1.1.2, Custom ResNet50, HT on, Turbo on, Comparing inference latency performance on same system with FP32 vs INT8 w/ Intel® DL Boost, test by Cloudwalk/Intel on 2/15/2019.

2.19x face recognition performance improvement for HiSign: Tested by Intel and HiSign as of 02/01/2019. 2 socket Intel® Xeon® Platinum 8260 Processor, 24 cores HT On Turbo ON Total Memory 768 GB (12 slots/ 64GB/ 2666 MHz), BIOS version 1.018 (ucode 0x400000A), RedHat 7.5 kernel 4.19.3-1.el7.elrepo.x86_64, Compiler: gcc 4.8.5, Deep Learning Framework: Intel® Optimizations for Caffe v1.1.2, Topology: modified Resnet32,custom dataset, BS=1. Comparing performance on same system with FP32 vs INT8 w/ Intel® DL Boost

2x Nokia* SDWAN: Configuration #1 (With Intel® QuickAssist® Technology): 2x Intel® Xeon® Gold 5218N Processor on Neon City Platform with 192 GB total memory (12 slots / 16GB / DDR4 2667MHz), ucode 0x4000019, Bios: PLYXCRB 1.86B.0568.D10.1901032132, uCode: 0x4000019 on CentOS 7.5 with Kernel 3.10.0-862, KVM Hypervisor; 1x Intel® QuickAssist Adapter 8970, Cipher: AES-128 SHA-256; Intel® Ethernet Converged Network Adapter X520-SR2; Application: Nokia Nuage SDWAN NSGv 5.3.3U3. Configuration # 2: 2x Intel® Xeon® Gold 5118 Processor on Neon City Platform with 192 GB total memory (12 slots / 16GB / DDR4 2667MHz), ucode 0x4000019, Bios: PLYXCRB 1.86B.0568.D10.1901032132, uCode: 0x4000019 on CentOS 7.5 with Kernel 3.10.0-862, KVM Hypervisor; Intel® Ethernet Converged Network Adapter X520-SR2; Application: Nokia Nuage SDWAN NSGv 5.3.3U3. Results recorded by Intel on 2/14/2018 in collaborate with Nokia.

SYSTEM CONFIGURATION: INTEL® DEEP LEARNING BOOST

1x inference throughput baseline on Intel® Xeon® Platinum 8180 processor (July 2017) : Tested by Intel as of July 11th 2017: Platform: 2S Intel® Xeon® Platinum 8180 CPU @ 2.50GHz (28 cores), HT disabled, turbo disabled, scaling governor set to "performance" via intel_pstate driver, 384GB DDR4-2666 ECC RAM. CentOS Linux release 7.3.1611 (Core), Linux kernel 3.10.0-514.10.2.el7.x86_64. SSD: Intel® SSD DC S3700 Series (800GB, 2.5in SATA 6Gb/s, 25nm, MLC).**Performance measured with**: Environment variables: KMP_AFFINITY='granularity=fine, compact', OMP_NUM_THREADS=56, CPU Freq set with cpupower frequency-set -d 2.5G -u 3.8G -g performance. Caffe: (http://github.com/intel/caffe/), revision f96b759f71b2281835f690af267158b82b150b5c. Inference measured with "caffe time --forward_only" command, training measured with "caffe time" command. For "ConvNet" topologies, synthetic dataset was used. For other topologies, data was stored on local storage and cached in memory before training. Topology specs from https://github.com/intel/caffe/tree/master/models/intel_optimized_models (ResNet-50),and https://github.com/soumith/convnet-benchmarks/tree/master/caffe/imagenet_winners (ConvNet benchmarks; files were updated to use newer Caffe prototxt format but are functionally equivalent). Intel C++ compiler ver. 17.0.2 20170213, Intel MKL small libraries version 2018.0.20170425. Caffe run with "numactl -l".

5.7x inference throughput improvement on Intel® Xeon® Platinum 8180 processor (December 2018) with continued optimizations : Tested by Intel as of November 11th 2018 :2 socket Intel(R) Xeon(R) Platinum 8180 CPU @ 2.50GHz / 28 cores HT ON , Turbo ON Total Memory 376.46GB (12slots / 32 GB / 2666 MHz). CentOS Linux-7.3.1611-Core, kernel: 3.10.0-862.3.3.el7.x86_64, SSD sda RS3WC080 HDD 744.1GB,sdb RS3WC080 HDD 1.5TB,sdc RS3WC080 HDD 5.5TB , Deep Learning Framework Intel® Optimization for caffe version: 551a53d63a6183c233abaa1a19458a25b672ad41 Topology::ResNet_50_v1 BIOS:SE5C620.86B.00.01.0014.070920180847 MKLDNN: 4e333787e0d66a1dca1218e99a891d493dbc8ef1 instances: 2 instances socket:2 (Results on Intel® Xeon® Scalable Processor were measured running multiple instances of the framework. Methodology described here: https://software.intel.com/en-us/articles/boosting-deep-learning-training-inference-performance-on-xeon-and-xeon-phi) Synthetic data. Datatype: INT8 Batchsize=64 vs Tested by Intel as of July 11th 2017:2S Intel® Xeon® Platinum 8180 CPU @ 2.50GHz (28 cores), HT disabled, turbo disabled, scaling governor set to "performance" via intel_pstate driver, 384GB DDR4-2666 ECC RAM. CentOS Linux release 7.3.1611 (Core), Linux kernel 3.10.0-514.10.2.el7.x86_64. SSD: Intel® SSD DC S3700 Series (800GB, 2.5in SATA 6Gb/s, 25nm, MLC).**Performance measured with**: Environment variables: KMP_AFFINITY='granularity=fine, compact', OMP_NUM_THREADS=56, CPU Freq set with cpupower frequency-set -d 2.5G -u 3.8G -g performance. Caffe: (http://github.com/intel/caffe/), revision f96b759f71b2281835f690af267158b82b150b5c. Inference measured with "caffe time --forward_only" command, training measured with "caffe time" command. For "ConvNet" topologies, synthetic dataset was used. For other topologies, data was stored on local storage and cached in memory before training. Topology specs from https://github.com/intel/caffe/tree/master/models/intel_optimized_models (ResNet-50). Intel C++ compiler ver. 17.0.2 20170213, Intel MKL small libraries version 2018.0.20170425. Caffe run with "numactl -l".

14x inference throughput improvement on Intel® Xeon® Platinum 8280 processor with Intel® DL Boost: Tested by Intel as of 2/20/2019. 2 socket Intel® Xeon® Platinum 8280 Processor, 28 cores HT On Turbo ON Total Memory 384 GB (12 slots/ 32GB/ 2933 MHz), BIOS: SE5C620.86B.0D.01.0271.120720180605 (ucode: 0x200004d), Ubuntu 18.04.1 LTS, kernel 4.15.0-45-generic, SSD 1x sda INTEL SSDSC2BA80 SSD 745.2GB, nvme1n1 INTEL SSDPE2KX040T7 SSD 3.7TB, Deep Learning Framework: Intel® Optimization for Caffe version: 1.1.3 (commit hash: 7010334f159da247db3fe3a9d96a3116ca06b09a) , ICC version 18.0.1, MKL DNN version: v0.17 (commit hash: 830a10059a018cd2634d94195140cf2d8790a75a, model: https://github.com/intel/caffe/blob/master/models/intel_optimized_models/int8/resnet50_int8_full_conv.prototxt, BS=64, synthetic Data, 4 instance/2 socket, Datatype: INT8 vs Tested by Intel as of July 11th 2017: 2S Intel® Xeon® Platinum 8180 CPU @ 2.50GHz (28 cores), HT disabled, turbo disabled, scaling governor set to "performance" via intel_pstate driver, 384GB DDR4-2666 ECC RAM. CentOS Linux release 7.3.1611 (Core), Linux kernel 3.10.0-514.10.2.el7.x86_64. SSD: Intel® SSD DC S3700 Series (800GB, 2.5in SATA 6Gb/s, 25nm, MLC).**Performance measured with**: Environment variables: KMP_AFFINITY='granularity=fine, compact', OMP_NUM_THREADS=56, CPU Freq set with cpupower frequency-set -d 2.5G -u 3.8G -g performance. Caffe: (http://github.com/intel/caffe/), revision f96b759f71b2281835f690af267158b82b150b5c. Inference measured with "caffe time --forward_only" command, training measured with "caffe time" command. For "ConvNet" topologies, synthetic dataset was used. For other topologies, data was stored on local storage and cached in memory before training. Topology specs from https://github.com/intel/caffe/tree/master/models/intel_optimized_models (ResNet-50),. Intel C++ compiler ver. 17.0.2 20170213, Intel MKL small libraries version 2018.0.20170425. Caffe run with "numactl -l".

2x More inference throughput improvement on Intel® Xeon® Platinum 9282 processor with Intel® DL Boost : Tested by Intel as of 2/26/2019. Platform: Dragon rock 2 socket Intel® Xeon® Platinum 9282(56 cores per socket), HT ON, turbo ON, Total Memory 768 GB (24 slots/ 32 GB/ 2933 MHz), BIOS:SE5C620.86B.0D.01.0241.112020180249, Centos 7 Kernel 3.10.0-957.5.1.el7.x86_64, Deep Learning Framework: Intel® Optimization for Caffe version: https://github.com/intel/caffe d554cbf1, ICC 2019.2.187, MKL DNN version: v0.17 (commit hash: 830a10059a018cd2634d94195140cf2d8790a75a), model: https://github.com/intel/caffe/blob/master/models/intel_optimized_models/int8/resnet50_int8_full_conv.prototxt, BS=64, No datalayer syntheticData:3x224x224, 56 instance/2 socket, Datatype: INT8 vs Tested by Intel as of July 11th 2017: 2S Intel® Xeon® Platinum 8180 CPU @ 2.50GHz (28 cores), HT disabled, turbo disabled, scaling governor set to "performance" via intel_pstate driver, 384GB DDR4-2666 ECC RAM. CentOS Linux release 7.3.1611 (Core), Linux kernel 3.10.0-514.10.2.el7.x86_64. SSD: Intel® SSD DC S3700 Series (800GB, 2.5in SATA 6Gb/s, 25nm, MLC).**Performance measured with**: Environment variables: KMP_AFFINITY='granularity=fine, compact', OMP_NUM_THREADS=56, CPU Freq set with cpupower frequency-set -d 2.5G -u 3.8G -g performance. Caffe: (http://github.com/intel/caffe/), revision f96b759f71b2281835f690af267158b82b150b5c. Inference measured with "caffe time --forward_only" command, training measured with "caffe time" command. For "ConvNet" topologies, synthetic dataset was used. For other topologies, data was stored on local storage and cached in memory before training. Topology specs from https://github.com/intel/caffe/tree/master/models/intel_optimized_models (ResNet-50),. Intel C++ compiler ver. 17.0.2 20170213, Intel MKL small libraries version 2018.0.20170425. Caffe run with "numactl -l".

SYSTEM CONFIGURATION: WORLD RECORD + REAL WORKLOAD PERFORMANCE LEADERSHIP

Performance results are based on testing as of dates shown in configuration and may not reflect all publicly available security updates. See configuration disclosure for details. No product or component 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 visit www.intel.com/benchmarks.

3.4X Facial Recognition for Microsoft: Tested by Intel as of 3/12/2019. Intel® Xeon® Platinum 8268 Processor, 24 cores, 384 GB (12 slots/ 32GB/ 2666 MHz), HT ON, BIOS: SE5C620.86B.BR.2018.6.10.1757, Ubuntu 4.19.5, 4.19.5-041905, nGraph version: b8106133dca9c63bf167e34306513111adf61995, ONNX version: 1.3.0, MKL DNN version: v0.18, MKLML_VERSION_2019.0.3.20190125,Topology: ResNet-50,BS=1, Dataset: Synthetic, Datatype: INT8 w/ Intel® DL Boost vs Tested by Intel as of 3/12/2019. Intel® Xeon® Platinum 8168 Processor, 24 cores, 384 GB (12 slots/ 32GB/ 2666 MHz), HT ON, BIOS: SE5C620.86B.BR.2018.6.10.1757, Ubuntu 4.19.5, 4.19.5-041905, nGraph version: b8106133dca9c63bf167e34306513111adf61995, ONNX version: 1.3.0, MKL DNN version: v0.18, MKLML_VERSION_2019.0.3.20190125,Topology: ResNet-50,BS=1, Dataset: Synthetic, Datatype: FP32

2.4x text detection performance improvement for JD.com: Tested by JD.com as of 1/27/2019. 2 socket Intel® Xeon® Gold Processor, 24 cores HT On Turbo ON Total Memory 192 GB (12 slots/ 16GB/ 2666 MHz), CentOS 7.6 3.10.0-957.el7.x86_64, Compiler: gcc 4.8.5, Deep Learning Framework: Intel® Optimizations for Caffe with custom optimizations, Topology: EAST (https://arxiv.org/abs/1704.03155), JD.com's private dataset, BS=1. Comparing performance on same system with FP32 vs INT8 w/ Intel® DL Boost

2.01x medical image classification performance improvement for NeuSoft: Tested by Intel and HiSign as of 02/01/2019. 2 socket Intel® Xeon® Platinum 8260 Processor, 24 cores HT On Turbo ON Total Memory 768 GB (12 slots/ 64GB/ 2666 MHz), BIOS version 1.018 (ucode 0x400000A), RedHat 7.5 kernel 4.19.3-1.el7.elrepo.x86_64, Compiler: gcc 4.8.5, Deep Learning Framework: Intel® Optimizations for Caffe v1.1.2, Topology: modified Alexnet ,custom dataset, BS=1. Comparing performance on same system with FP32 vs INT8 w/ Intel® DL Boost

4.43X ML Inferencing for Target: Based on Intel Analysis on 2/16/2019. 2nd Gen Intel® Xeon® Platinum 8280 Processor (28 Cores) with 384GB, DDR4-2933, using Intel® OpenVino™ 2019 R1. HT OFF, Turbo ON. CentOS Linux release 7.6.1810, kernel 4.19.5-1.el7.elrepo.x86_64. Topology: ResNet-50, dataset: Synthetic, BS=4 and 14 instance, Comparing FP32 vs Int8 w/ Intel® DL Boost performance on the system.

3.26x latency reduction for Tencent* Cloud Video Analysis: Tested by Tencent as of 1/14/2019. 2 socket Intel® Xeon® Gold Processor, 24 cores HT On Turbo ON Total Memory 192 GB (12 slots/ 16GB/ 2666 MHz), CentOS 7.6 3.10.0-957.el7.x86_64, Compiler: gcc 4.8.5, Deep Learning Framework: Intel® Optimizations for Caffe v1.1.3, Topology: modified inception v3, Tencent's private dataset, BS=1. Comparing performance on same system with FP32 vs INT8 w/ Intel® DL Boost

SYSTEM CONFIGURATION: SKUS OPTIMIZED FOR UNIQUE NETWORK NEEDS

Performance results are based on testing as of dates shown in configuration and may not reflect all publicly available security updates. See configuration disclosure for details. No product or component 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 visit www.intel.com/benchmarks.

Up to 1.76x gains on networking workloads based on OVS DPDK: Tested by Intel on 1/21/2019 1-Node, 2x Intel® Xeon® Gold 6130 Processor on Neon City platform with 12x 16GB DDR4 2666MHz (384GB total memory), Storage: 1x Intel® 240GB SSD, Network: 4x Intel XXV710-DA2, Bios: PLYXCRB1.86B.0568.D10.1901032132, ucode: 0x200004d (HT= ON, Turbo= OFF), OS: Ubuntu* 18.04 with kernel: 4.15.0-42-generic, Benchmark: Open Virtual Switch (on 4C/4P/8T 64B Mpacket/s), Workload version: OVS 2.10.1, DPDK-17.11.4, Compiler: gcc7.3.0, Other software: QEMU-2.12.1, VPP v18.10, Results: 9.6. Tested by Intel on 1/18/2019 1-Node, 2x Intel® Xeon® Gold 6230N Processor on Neon City platform with 12x 16GB DDR4 2999MHz (384GB total memory), Storage: 1x Intel® 240GB SSD, Network: 6x Intel XXV710-DA2, Bios: PLYXCRB1.86B.0568.D10.1901032132, ucode: 0x4000019 (HT= ON, Turbo= OFF), OS: Ubuntu* 18.04 with kernel: 4.20.0-042000rc6-generic, Benchmark: Open Virtual Switch (on 6P/6C/12T 64B Mpacket/s), Workload version: OVS 2.10.1, DPDK-17.11.4, Compiler: gcc7.3.0, Other software: QEMU-2.12.1, VPP v18.10, Results: 15.2. Tested by Intel on 1/18/2019 1-Node, 2x Intel® Xeon® Gold 6230N Processor with SST-BF enabled on Neon City platform with 12x 16GB DDR4 2999MHz (384GB total memory), Storage: 1x Intel® 240GB SSD, Network: 6x Intel XXV710-DA2, Bios: PLYXCRB1.86B.0568.D10.1901032132, ucode: 0x4000019 (HT= ON, Turbo= ON (SST-BF)), OS: Ubuntu* 18.04 with kernel: 4.20.0-042000rc6-generic, Benchmark: Open Virtual Switch (on 6P/6C/12T 64B Mpacket/s), Workload version: OVS 2.10.1, DPDK-17.11.4, Compiler: gcc7.3.0, Other software: QEMU-2.12.1, VPP v18.10, Results: 16.9