intel

# DATA-CENTRIC INNOVATION SUMMIT

AUGUST 8, 2018 | SANTA CLARA, CA

DATA-CENTRIC INNOVATION SUMMIT

# NAVIN SHENOY

EXECUTIVE VICE PRESIDENT & GENERAL MANAGER
DATA CENTER GROUP

# DATA DEFINES THE FUTURE

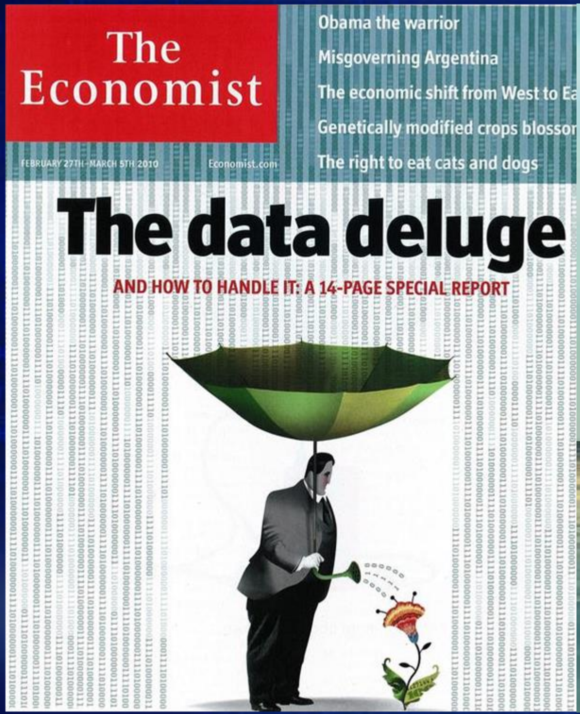




# DATA-CENTRIC TRANSFORMATION

## DECREASING COST OF TECHNOLOGY

**COST OF COMPUTE**

**56%** 2012-2017

**COST OF STORAGE**

**77%** 2012-2017

Source: Amalgamation of analyst data and Intel analysis.

## 41X PERFORMANCE INCREASE

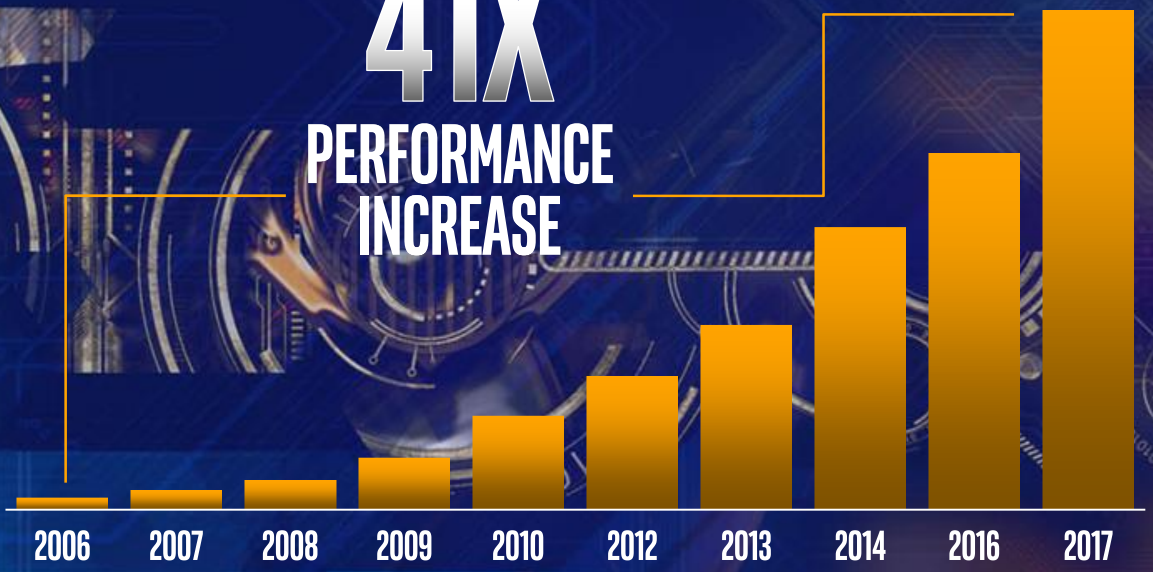


Source: Intel

# GREATEST DATA COLLECTOR

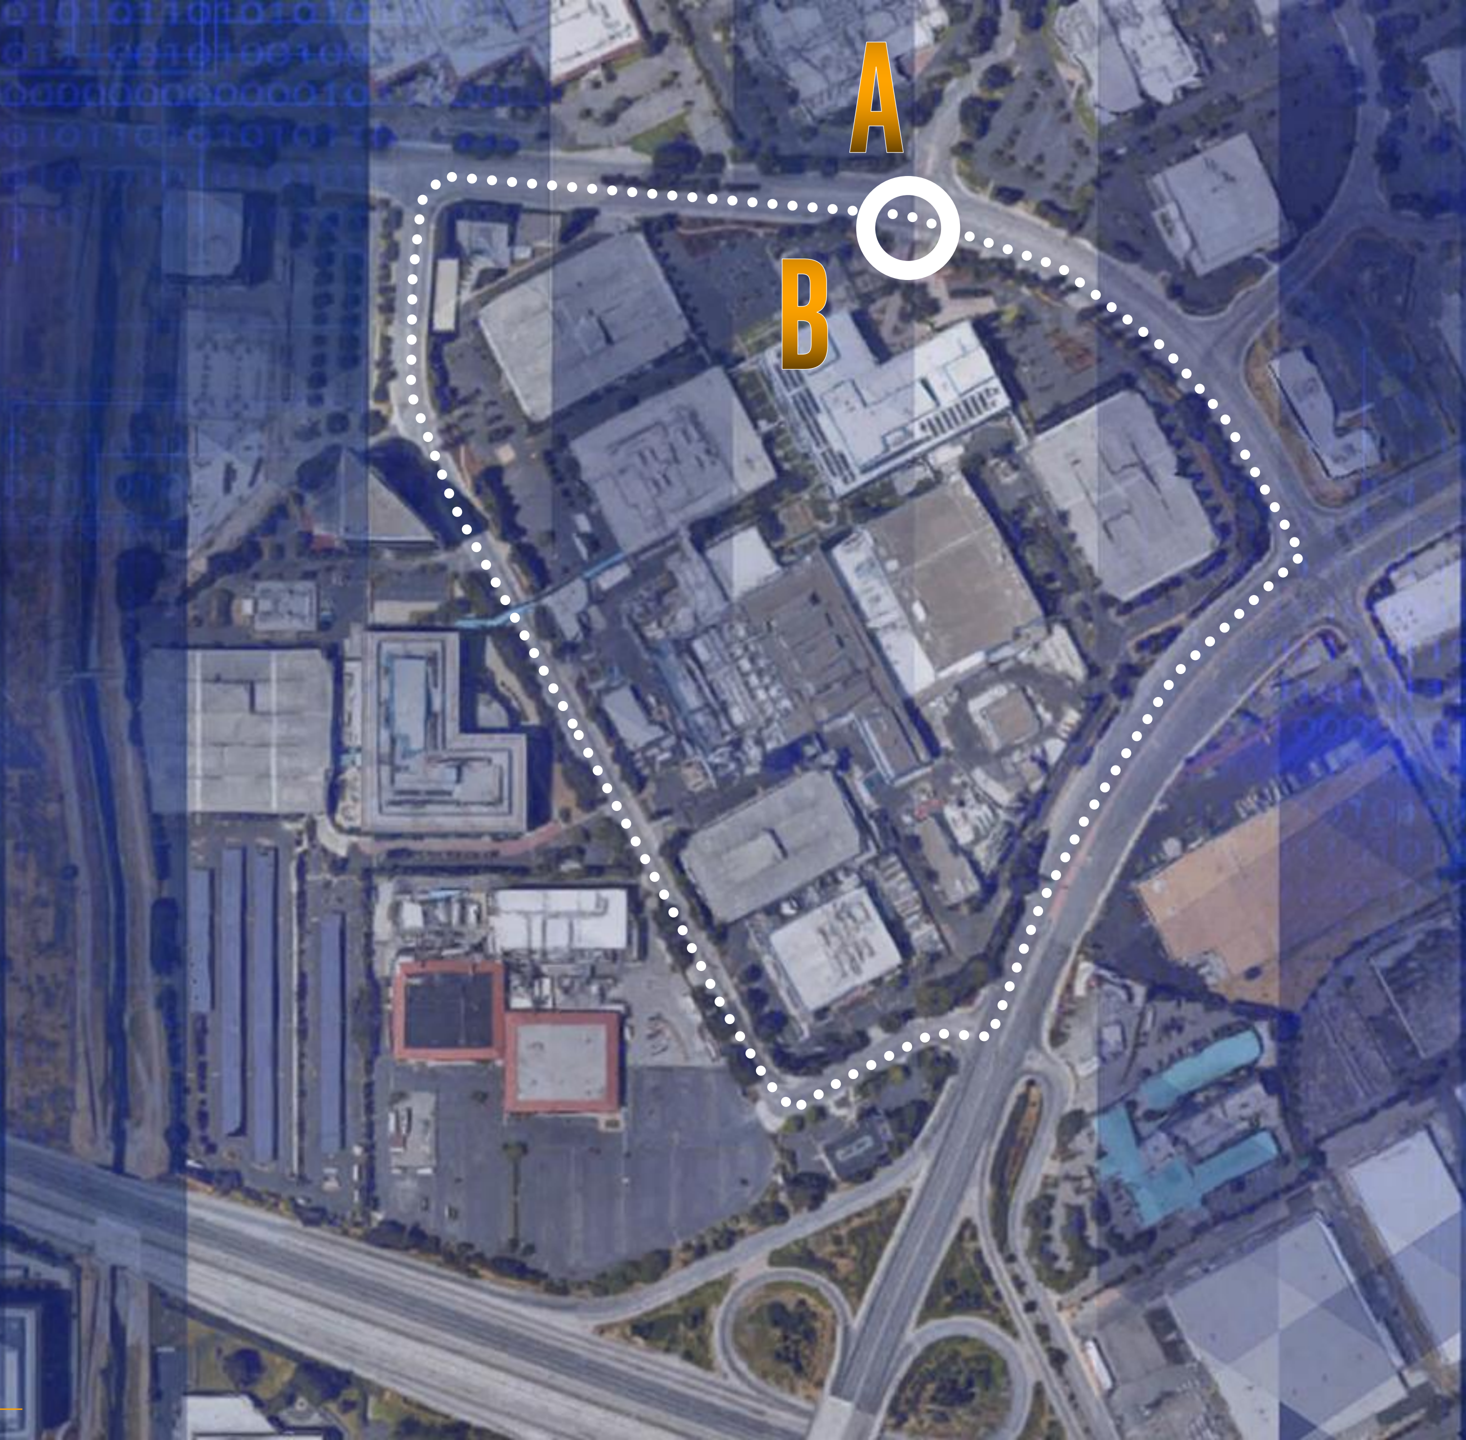


DATA-CENTRIC
INNOVATION SUMMIT

intel


FRONT FACING CAMERAS
REAR FACING CAMERA
LATERAL CAMERA


DATA-CENTRIC
INNOVATION SUMMIT

DATA-CENTRIC
INNOVATION SUMMIT


EDGE
CLOUD
AI
VISION
DATA-CENTRIC
INNOVATION SUMMIT

# DATA-CENTRIC OPPORTUNITY

## 2021 DATA-CENTRIC SI TAM

FROM 2017

| ~$70B | ~$55B | ~$30B | ~$7B |
|---|---|---|---|
| NETWORK<br>MEMORY<br>CONNECTIVITY<br>AI | INTEL OPTANE SSDS<br>3D NAND | ADAS<br>INDUSTRIAL<br>VIDEO<br>RETAIL<br>AI | ADAS<br>NETWORK<br>DATA CENTER<br>AI |
| 2017 REVENUE | | | |
| DATA CENTER | NON VOLATILE MEMORY | IOT + ADAS | FPGA |

**TOTAL TAM >$160B**

Source: 2017 Intel Revenue is based on Intel financials. 2021F 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. Data Center includes Server, Storage, & Network computing, Ethernet/OPA, Silicon Photonics and Memory. Non-Volatile Memory includes NAND and 3D XPoint™ technology. IOT includes addressable Logic ASIC/ASSP, MPU, MCU, DSP for Industrial, Transportation, Automated Driving, Retail, Video Surveillance, Healthcare, Public Sector, Office Automation, Gaming and Smart Home. PSG includes Field Programmable Logic.
Forecast is soft estimate subject to significant change and revisions

# DATA-CENTRIC OPPORTUNITY

## 2022 DATA-CENTRIC SI TAM

| DATA CENTER | NON VOLATILE MEMORY | IOT + ADAS | FPGA |
|---|---|---|---|
| ~$90B | ~$75B | ~$33B | ~$8B |
| ~~~$70B~~ | ~~~$55B~~ | ~~~$30B~~ | ~~~$7B~~ |
| NETWORK<br>MEMORY<br>CONNECTIVITY<br>AI | INTEL OPTANE SSDS<br>3D NAND | ADAS<br>INDUSTRIAL<br>VIDEO<br>RETAIL<br>AI | ADAS<br>NETWORK<br>DATA CENTER<br>AI |
| 2017 REVENUE | | | |

**TOTAL TAM >$200B**

9% CAGR
2017-2022

~~TOTAL TAM >$160B~~

Source: 2017 Intel Revenue is based on Intel financials. 2021F 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. Data Center includes Server, Storage, & Network computing, Ethernet/OPA, Silicon Photonics and Memory. Non-Volatile Memory includes NAND and 3D XPoint™ technology. IOT includes addressable Logic ASIC/ASSP, MPU, MCU, DSP for Industrial, Transportation, Automated Driving, Retail, Video Surveillance, Healthcare, Public Sector, Office Automation, Gaming and Smart Home. PSG includes Field Programmable Logic.
Forecast is soft estimate subject to significant change and revisions

# CLOUD EXPANDING THE TAM

## PUBLIC | PRIVATE | HYBRID

INTEL CLOUD SP REVENUE

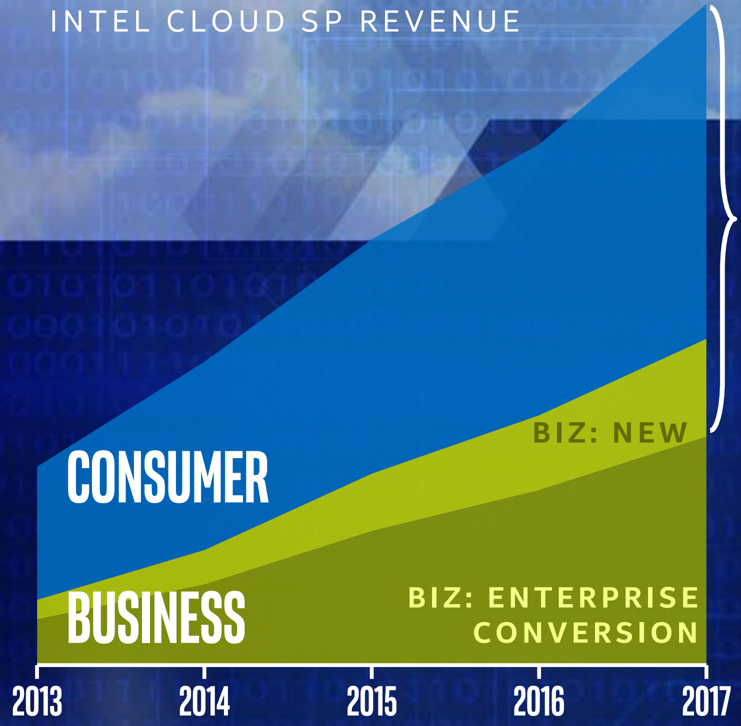


**2/3 OF CLOUD IS TAM EXPANSION**

## INCREASING NEED FOR CUSTOM CPUS

INTEL CLOUD SP CPU VOLUME

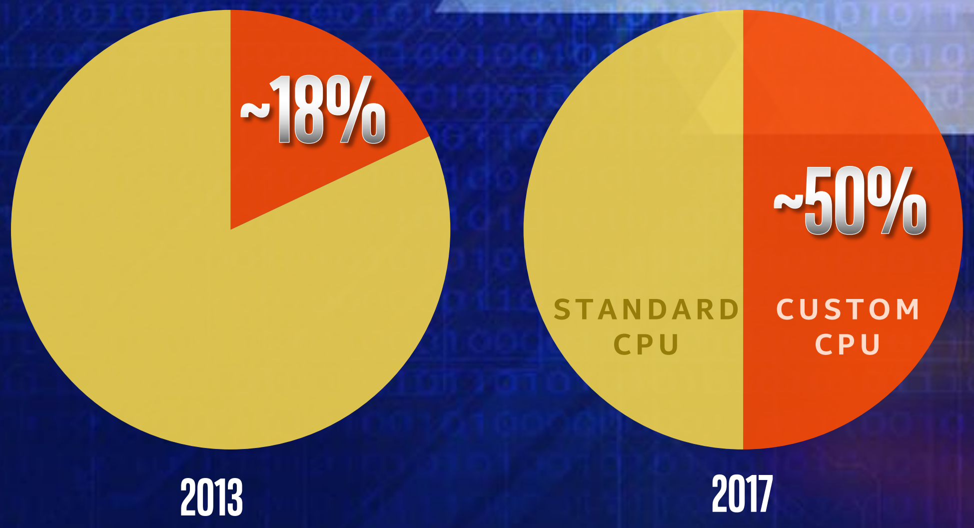




Source: Intel

# CLOUD-IFICATION OF
# NETWORK ◦ 5G ◦ EDGE

DEVICES | THINGS ACCESS | EDGE CORE DATA CENTER | CLOUD

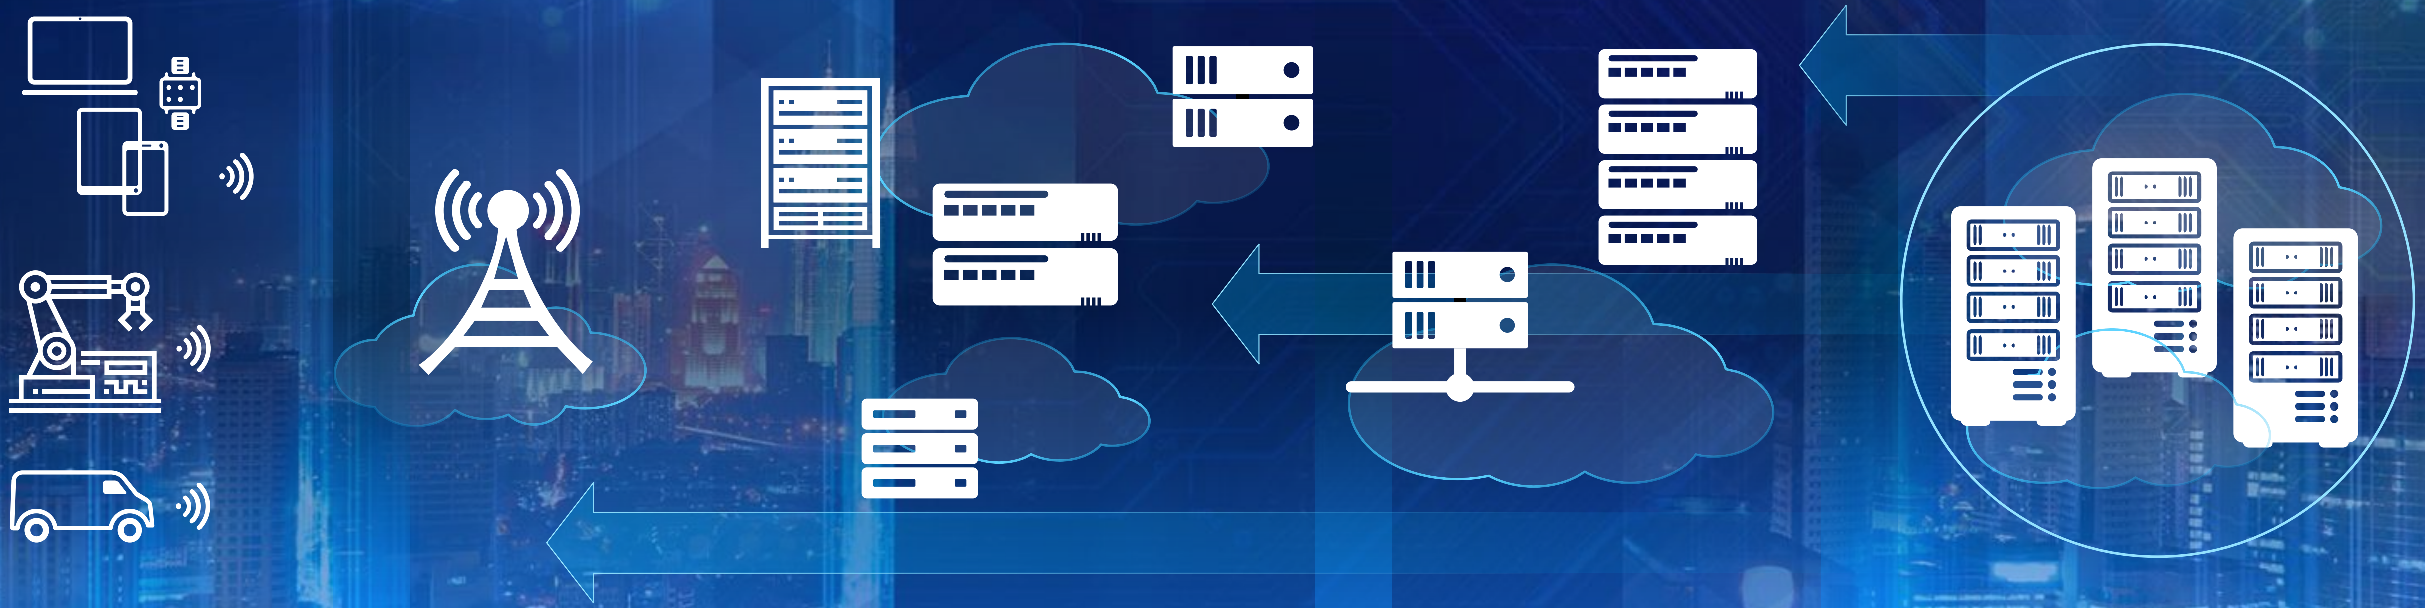

## $24B NETWORK LOGIC SILICON TAM | 2022

DATA-CENTRIC INNOVATION SUMMIT

Source: Si TAM projections and growth estimates are based on an amalgamation of analyst data and Intel analysis, reflect current expectations and available information, and are subject to change without notice. Network Si TAM includes CPUs, ASICs, FPGAs, SOCs, and other types of processors for network appliances

# ARTIFICIAL INTELLIGENCE

## AI DATA CENTER LOGIC SILICON TAM

~30% CAGR

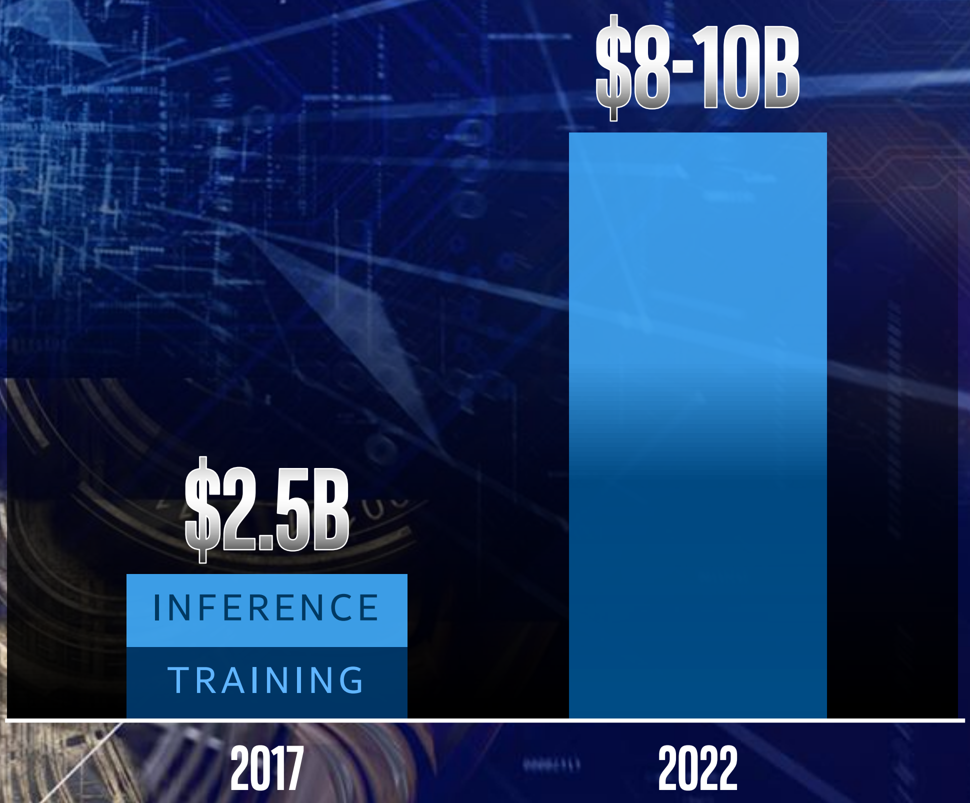




Source: AI Si TAM is based on amalgamation of analyst data and Intel analysis, based upon current expectations and available information and are subject to change without notice.

# NEW ERA OF DATA CENTER TECHNOLOGY

## DATA-CENTRIC INFRASTRUCTURE

### MOVE FASTER

intel SILICON PHOTONICS

intel OMNI-PATH FABRIC

intel STRATIX 10 inside

intel ETHERNET

### STORE MORE

intel OPTANE DC SOLID STATE DRIVE

intel OPTANE DC PERSISTENT MEMORY

### PROCESS EVERYTHING

intel XEON PLATINUM inside

intel STRATIX 10 inside

intel NERVANA inside

# DATA CENTER NETWORK TRAFFIC INCREASING

## GLOBAL DATA CENTER TRAFFIC PER YEAR

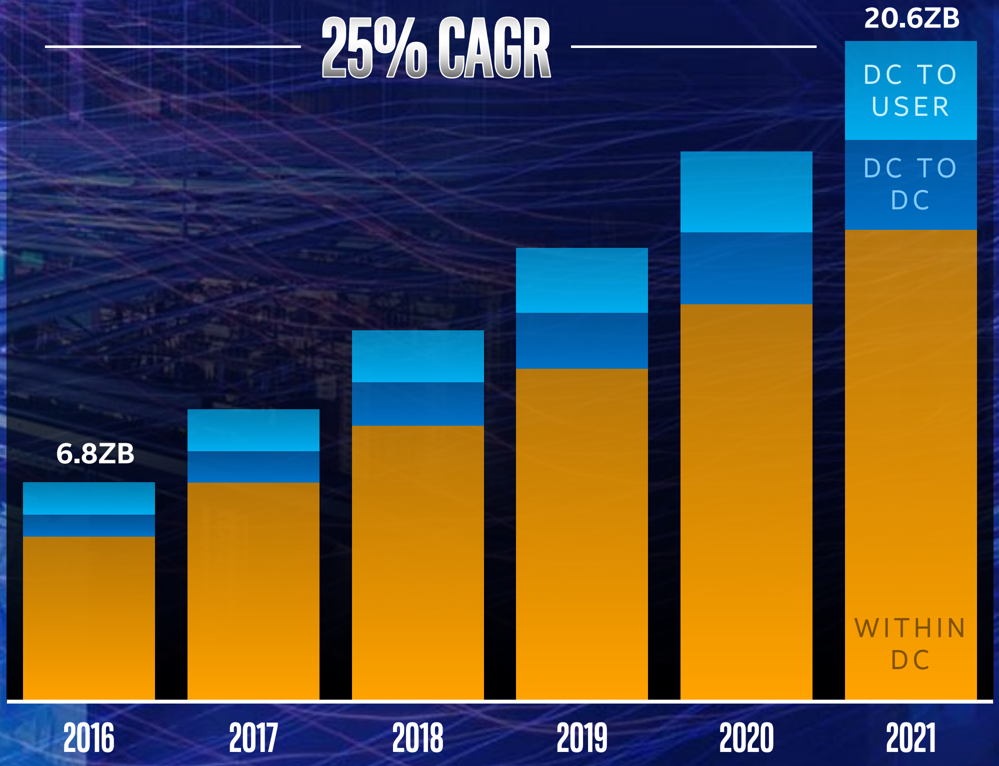




Source: Cisco Global Cloud Index, 2016-2021

# INTEL CONNECTIVITY PORTFOLIO

## CONNECTIVITY LOGIC SILICON TAM

~25% CAGR

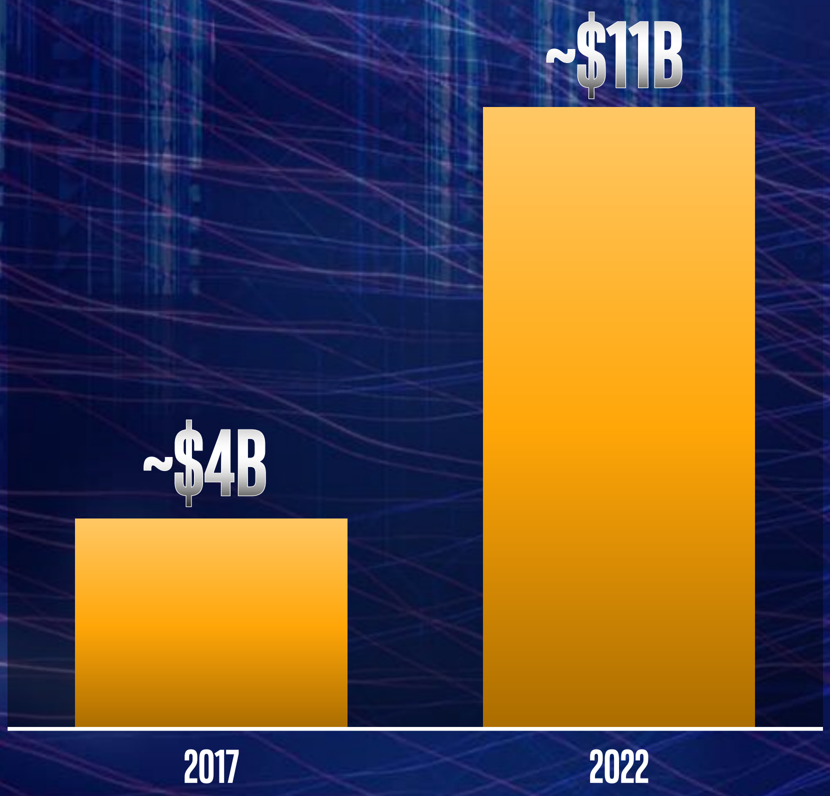


### INTEL® OMNI-PATH FABRIC

LEADING HPC FABRICS

### INTEL® ETHERNET

#1 MSS HIGH SPEED[1] ETHERNET

COMING 2019
CASCADE GLACIER SMARTNIC

### INTEL® SILICON PHOTONICS

SILICON INTEGRATION
SILICON MANUFACTURING
SILICON SCALE



1. High speed = 10GbE and above

Source: Connectivity TAM includes Ethernet, High Performance Fabrics, and Silicon Photonics and is based on amalgamation of analyst data and Intel analysis, based upon current expectations and available information and are subject to change without notice..

# NEW ERA OF DATA CENTER TECHNOLOGY

## DATA-CENTRIC INFRASTRUCTURE

### MOVE FASTER

intel SILICON PHOTONICS

intel OMNI-PATH FABRIC

intel STRATIX 10 inside

intel ETHERNET

### STORE MORE

intel OPTANE DC SOLID STATE DRIVE

intel OPTANE DC PERSISTENT MEMORY

### PROCESS EVERYTHING

intel XEON PLATINUM inside

intel STRATIX 10 inside

intel NERVANA inside

# RE-ARCHITECTING THE MEMORY / STORAGE HIERARCHY

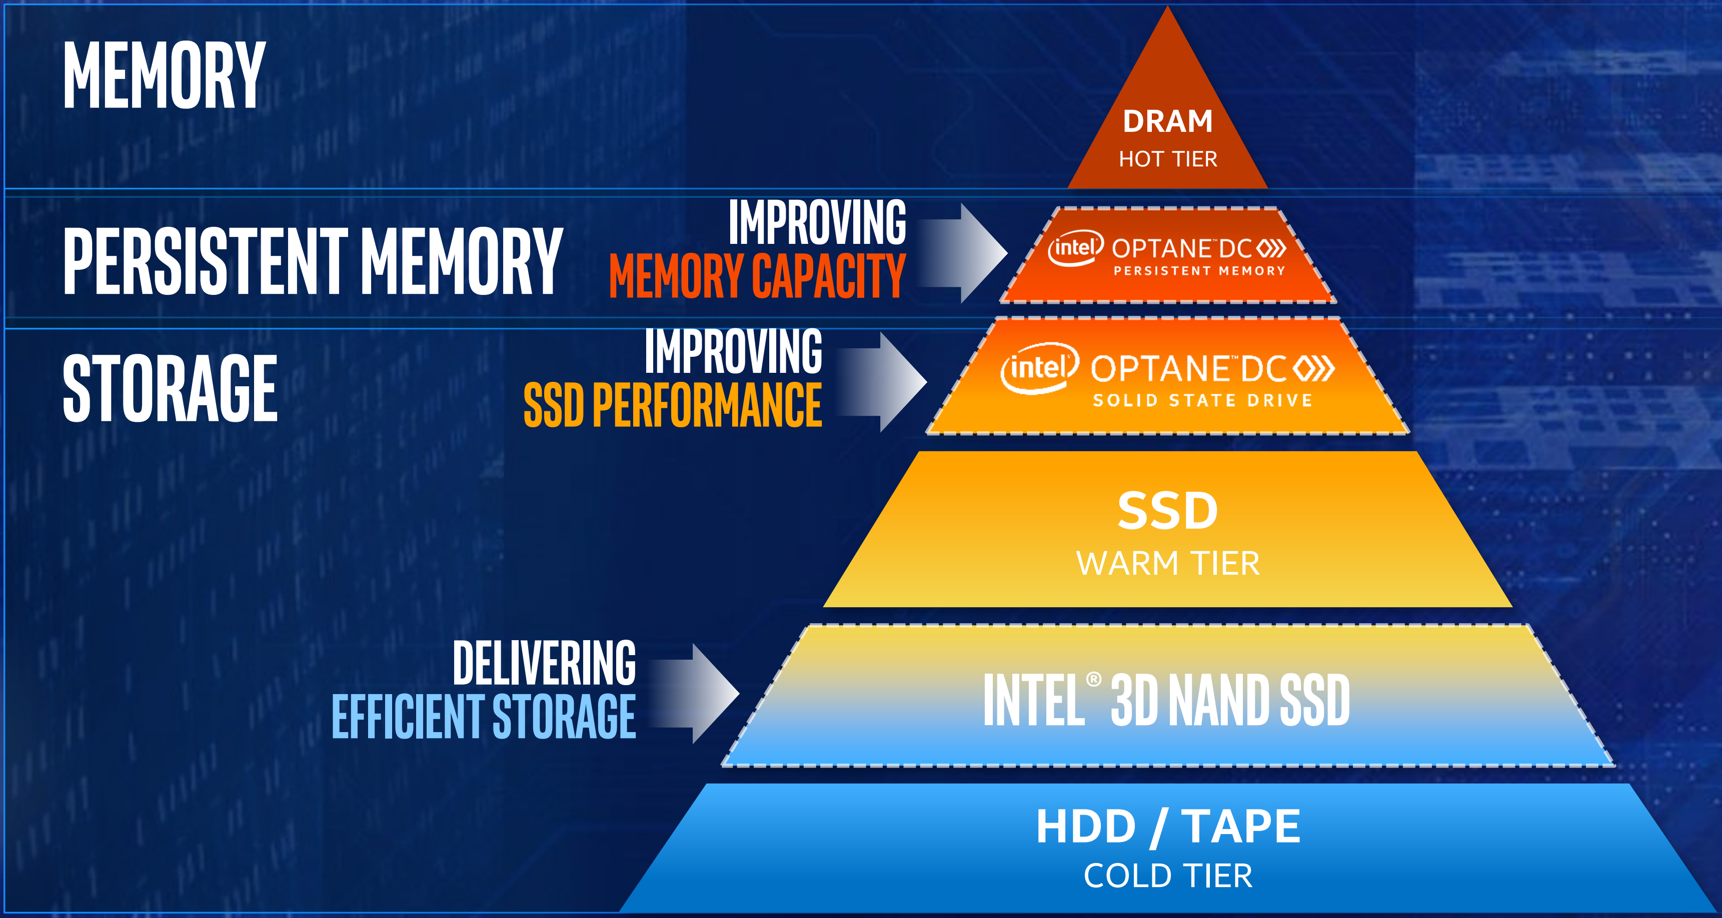

intel OPTANE™ DC
PERSISTENT MEMORY

# DATA CENTER MEMORY SAM | 2022

# $10B

## SPARK SQL DS

**8X** MORE PERFORMANCE

VS. DRAM AT 2.6TB DATA SCALE

## APACHE CASSANDRA

**9X** MORE READ TRANSACTIONS

**11X** MORE USERS PER SYSTEM

VS. COMPARABLE SERVER SYSTEM WITH DRAM & NAND NVME DRIVES

## VALUE OF PERSISTENCE

MINUTES TO **SECONDS**
START TIME

THREE 9s TO **FIVE 9s**
AVAILABILITY

## UNIQUE INTEL PLATFORM DIFFERENTIATION

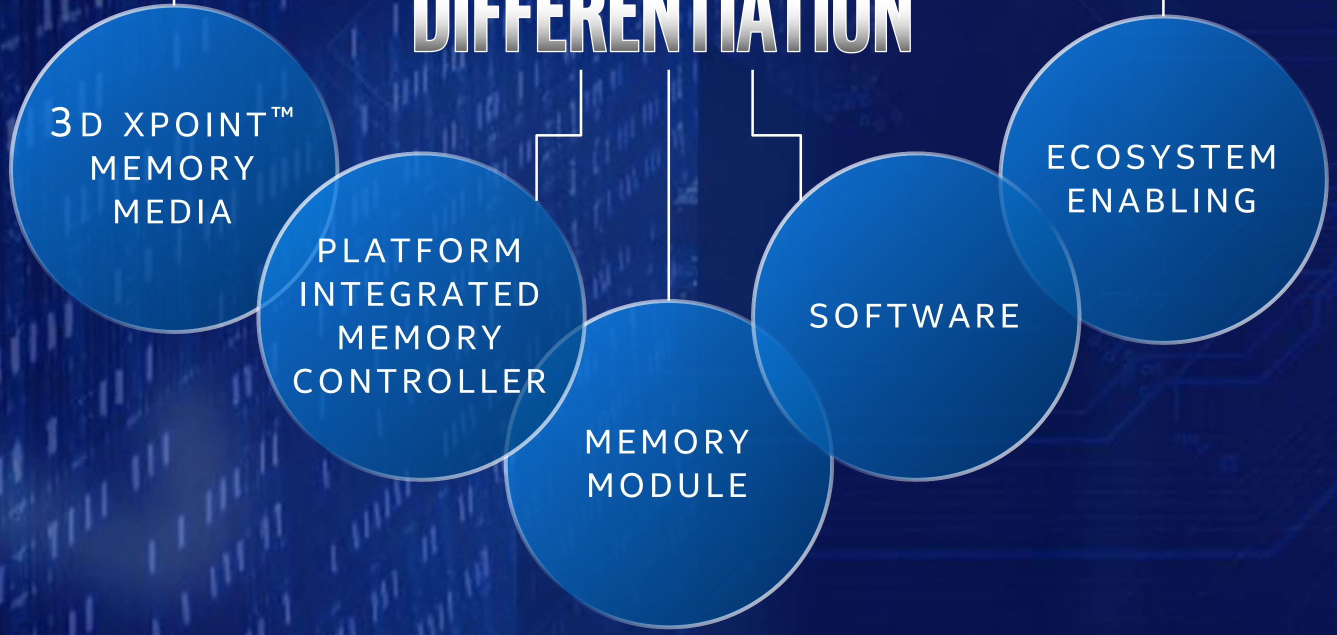


Note: Intel Optane DC Persistent Memory opportunity equals 2022 data center memory SAM. Performance results are based on testing: 8X (8/2/2018), 9X Reads/11X Users (5/24/2018), Minutes to Seconds (5/30/2018) and may not reflect all publicly available security updates. No product can be absolutely secure. See configuration disclosure for details. 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 information go to www.intel.com/benchmarks. Other names and brands may be claimed as the property of others.

# BART SANO

VP OF PLATFORMS

Google Cloud





intel

Google Cloud & intel

# BROAD ECOSYSTEM SUPPORT

AEROSPIKE
Alibaba.com
ALTIBASE
APACHE Spark
AsiaInfo
aws
cassandra
cisco
databricks
DELLEMC
FUJITSU
GIGASPACES innovate with confidence
Google Cloud
Hewlett Packard Enterprise
HUAWEI
IBM
人大金仓 Kingbase
金证股份
Lenovo
Microsoft
ORACLE
redhat
redislabs
RocksDB
SAP
sas
SUNJESOFT SUNJESOFT Inc.
SUSE
Tencent Cloud
天源迪科
ubuntu
vmware

# NEW ERA OF DATA CENTER TECHNOLOGY

## DATA-CENTRIC INFRASTRUCTURE

### MOVE FASTER

intel SILICON PHOTONICS

intel OMNI-PATH FABRIC

intel STRATIX 10 inside

intel ETHERNET

### STORE MORE

intel OPTANE DC SOLID STATE DRIVE

intel OPTANE DC PERSISTENT MEMORY

### PROCESS EVERYTHING

intel XEON PLATINUM inside

intel STRATIX 10 inside

intel NERVANA inside

# 20TH ANNIVERSARY OF THE INTEL® XEON® PROCESSOR

1998 | 1999 | 2000 | 2001 | 2002 | 2003 | 2004 | 2005 | 2006 | 2007 | 2008 | 2009 | 2010 | 2011 | 2012 | 2013 | 2014 | 2015 | 2016 | 2017 | 2018

SKYLAKE
BROADWELL
HASWELL
IVY BRIDGE
SANDY BRIDGE
WESTMERE | BECKTON
NEHALEM
ALLENDALE | WOLFDALE | DUNNINGTON
YORKFIELD | HARPERTOWN
DEMPSEY | SOSSAMAN | WOODCREST
IRWINDALE | PAXVILLE
NOCONA
PRESTONIA | GALLATIN
FOSTER
TANNER | CASCADES
DRAKE

intel® XEON® PLATINUM inside™

DATA-CENTRIC INNOVATION SUMMIT

intel

# 20TH ANNIVERSARY OF THE INTEL® XEON® PROCESSOR

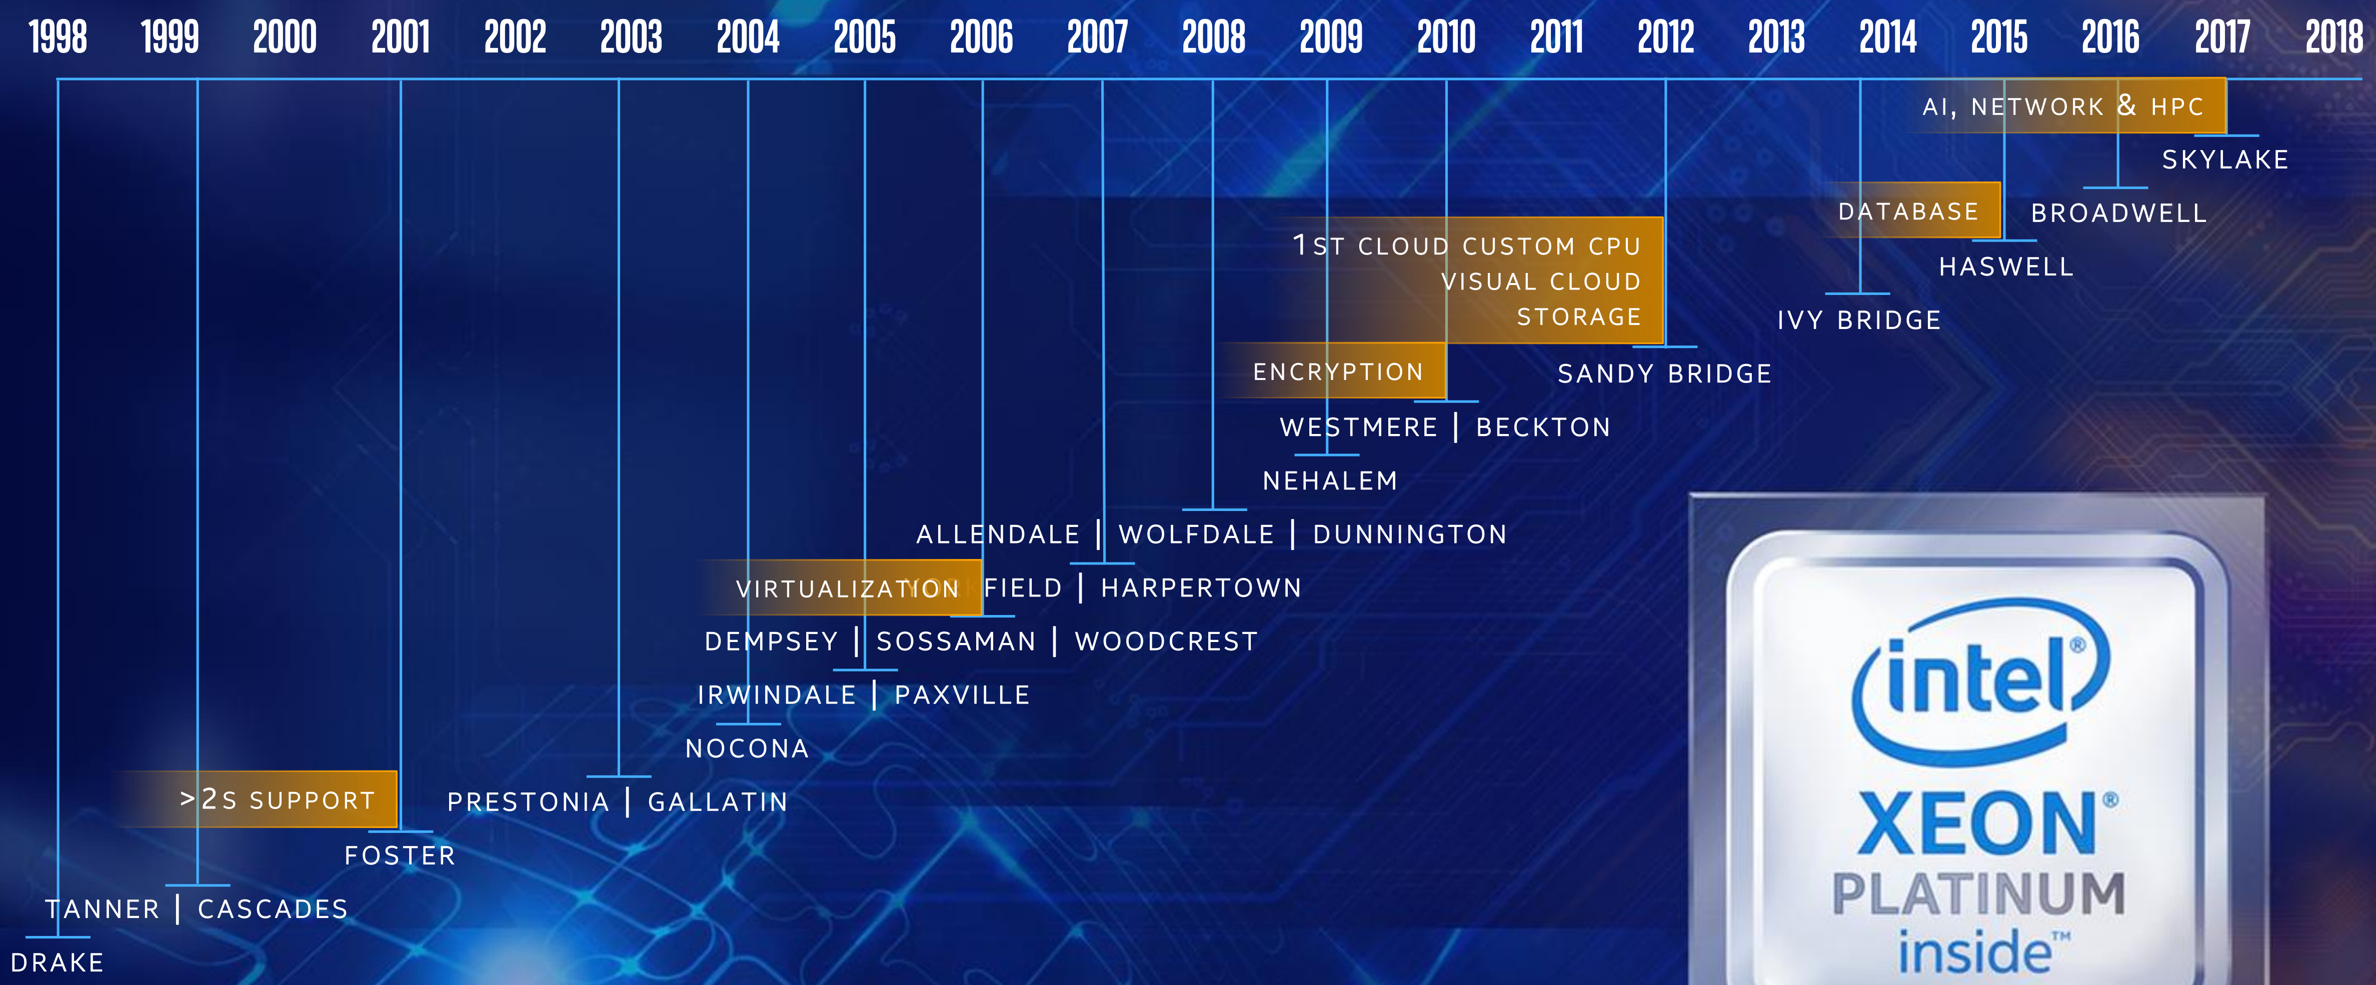


intel XEON® PLATINUM inside™

# 1ST ANNIVERSARY OF THE INTEL® XEON® SCALABLE PROCESSOR

**LARGEST** EARLY SHIP PROGRAM

**FASTEST** XEON RAMP TO 1M UNITS

**50%** OF XEON VOLUME

**>2M** UNITS SHIPPING PER QUARTER

## LEADERSHIP PERFORMANCE

VS OTHER x86 OFFERINGS

| | | | | |
|---|---|---|---|---|
| UP TO **1.48X** PER CORE | UP TO **1.72X** L3 PACKET FWD | UP TO **3.2X** HIGH PERF. LINPACK | UP TO **1.85X** DATABASE | UP TO **1.45X** MEMORY CACHING |

## ULTIMATE FLEXIBILITY

| | | | | |
|---|---|---|---|---|
| **1,2,4,8+** SOCKETS | **60** SKUS | **1.7-3.6** GHZ | **70-205** WATTS | **$213-$10,000** PRICE POINTS |

DATA-CENTRIC INNOVATION SUMMIT



intel

# REINVENTING XEON FOR AI

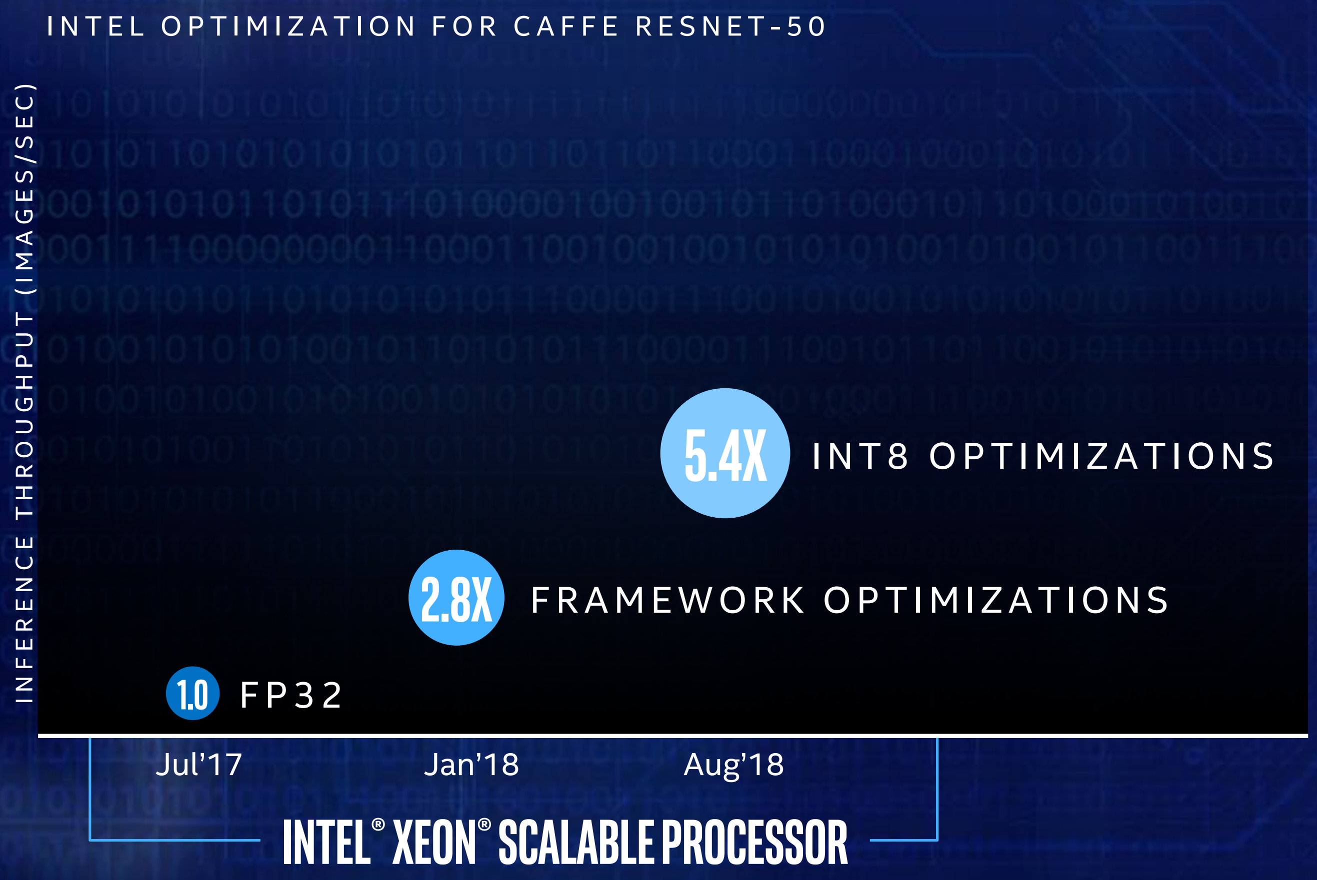


1 Intel® Optimization for Caffe Resnet-50 performance does not necessarily represent other Framework performance.
2 Based on Intel internal testing: 1X (7/11/2017), 2.8X (1/19/2018) and 5.4X (7/26/2018) performance improvement based on Intel® Optimization for Café Resnet-50 inference throughput performance on Intel® Xeon® Scalable Processor.
3 11X (7/25/2018) Results have been estimated using internal Intel analysis, and provided to you for informational purposes. Any differences in your system hardware, software or configuration may affect your actual performance.
Performance results are based on testing as of 7/11/2017(1x), 1/19/2018(2.8x) & 7/26/2018(5.4) and may not reflect all publically available security update. See configuration disclosure for details (config 1). 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.




intel

# WINNING AI ON IA

"Machine learning is a big part of our heritage. It works on GPUs today, but it also works on instances **powered by highly customized Intel Xeon processors**"

Bratin Saha
VP & GM, Machine Learning Platforms
Amazon AI - Amazon

"Inference is one thing we do, but we do lots more. That's why **flexibility is really essential.**"

Kim Hazelwood
Head of AI Infrastructure Foundation
Facebook

INTEL® XEON® PROCESSOR AI WINS

Alibaba Cloud
aws
bluedata
cdhi
CRAY
DELL EMC
GIGASPACES innovate with confidence
Google Cloud
科大讯飞 iFLYTEK
JD.COM 京东
KYOTO UNIVERSITY
MTC
NOVARTIS
NYU Langone MEDICAL CENTER
Taboola
头条 TOUTIAO
UCLOUD
UnionPay 银联

DATA-CENTRIC INNOVATION SUMMIT



intel

# WINNING AI ON IA

"Machine learning is a big part of our heritage. It works on GPUs today, but it also works on instances **powered by highly customized Intel Xeon processors**"

Bratin Saha
VP & GM, Machine Learning Platforms
Amazon AI - Amazon

## IN 2017 AI DROVE >$1B INTEL XEON REVENUE

"Inference is one thing we do, but we do lots more. That's why **flexibility is really essential.**"

Kim Hazelwood
Head of AI Infrastructure Foundation
Facebook

INTEL® XEON® PROCESSOR AI WINS

Alibaba Cloud, aws, bluedata, cdhi, CRAY, DELL EMC, GIGASPACES innovate with confidence, Google Cloud, 科大讯飞 iFLYTEK, JD.COM 京东, KYOTO UNIVERSITY, MTC, NOVARTIS, NYU Langone MEDICAL CENTER, Taboola, 头条 TOUTIAO, UCLOUD, UnionPay 银联

DATA-CENTRIC INNOVATION SUMMIT



intel

NEXT INTEL® XEON® SCALABLE PROCESSOR

# CASCADE LAKE

## WITH INTEL® OPTANE™ DC PERSISTENT MEMORY

Leadership Performance

Optimized Cache Hierarchy

Higher Frequencies

intel® XEON® PLATINUM inside™

Support For intel® OPTANE™ DC PERSISTENT MEMORY

Security Mitigations

Optimized Frameworks & Libraries

# REINVENTING XEON FOR AI

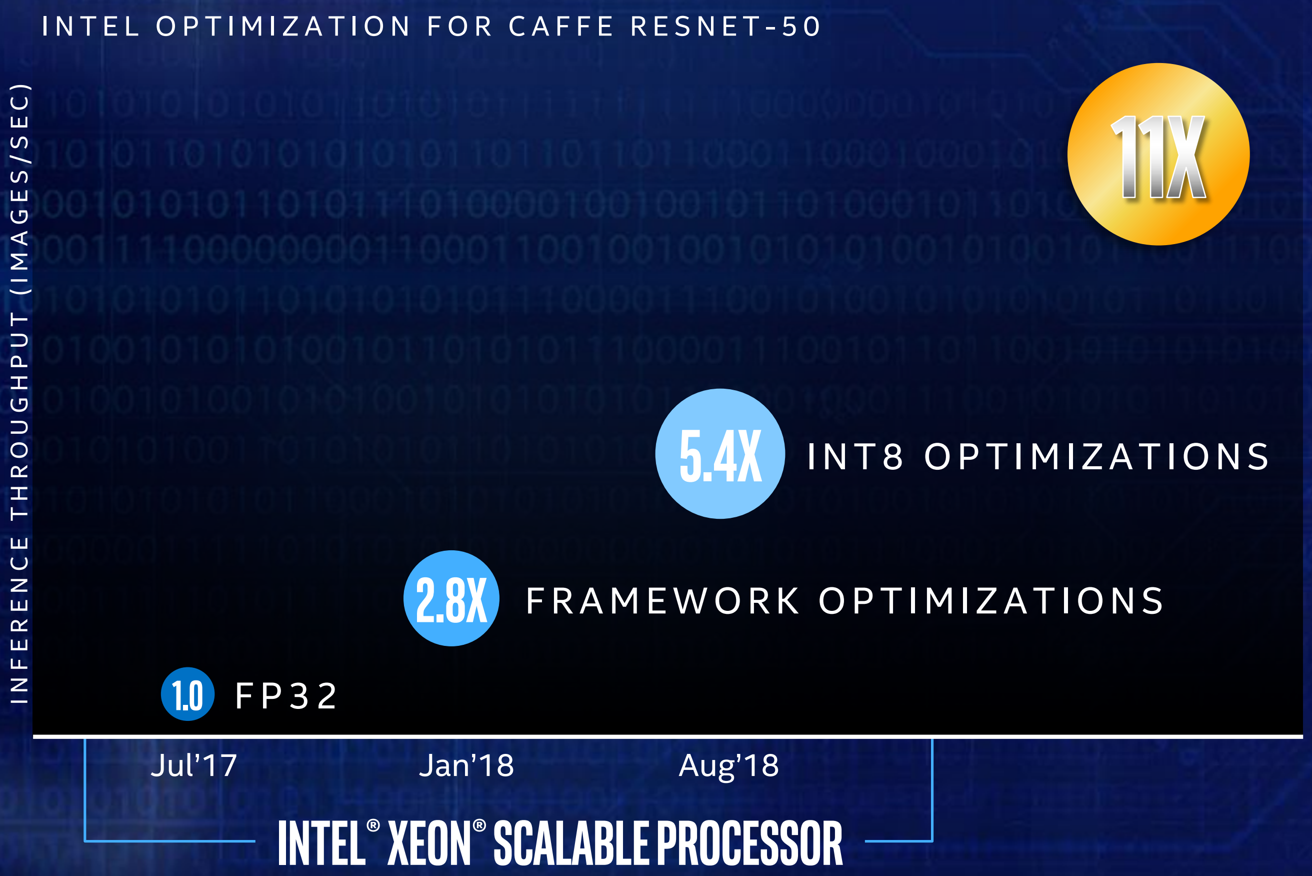


INTRODUCING

## INTEL® DEEP LEARNING BOOST

**VNNI** VECTOR NEURAL NETWORK INSTRUCTION

FOR INFERENCE ACCELERATION

### FRAMEWORK & LIBRARY SUPPORT

Caffe | mxnet | TensorFlow

intel MKL-DNN

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) and 5.4X (7/26/2018) performance improvement based on Intel® Optimization for Café Resnet-50 inference throughput performance on Intel® Xeon® Scalable Processor.
3 11X (7/25/2018) Results have been estimated using internal Intel analysis, and provided to you for informational purposes. Any differences in your system hardware, software or configuration may affect your actual performance.
Performance results are based on testing as of 7/11/2017(1x), 1/19/2018(2.8x) & 7/26/2018(5.4) and may not reflect all publically available security update. See configuration disclosure for details (config 1). 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.

# INTEL SELECT SOLUTIONS

intel select solution

FASTER. EASIER. OPTIMIZED.

TIGHTLY SPECIFIED HW & SW COMPONENTS — **SIMPLIFIED EVALUATION**

PRE-DEFINED SETTINGS & SYSTEM-WIDE TUNING — **FAST & EASY TO DEPLOY**

DESIGNED TO DELIVER OPTIMAL PERFORMANCE — **WORKLOAD OPTIMIZED**

Intel® Select Solution configurations and benchmark results are

**INTEL VERIFIED**

# INTEL SELECT SOLUTIONS

intel select solution
FASTER. EASIER. OPTIMIZED.

TIGHTLY SPECIFIED HW & SW COMPONENTS — **SIMPLIFIED EVALUATION**

PRE-DEFINED SETTINGS & SYSTEM-WIDE TUNING — **FAST & EASY TO DEPLOY**

DESIGNED TO DELIVER OPTIMAL PERFORMANCE — **WORKLOAD OPTIMIZED**

## INTRODUCING

INTEL SELECT SOLUTION
**AI: BIG DL ON APACHE SPARK**

INTEL SELECT SOLUTION
**BLOCKCHAIN: HYPERLEDGER FABRIC**

INTEL SELECT SOLUTION
**SAP HANA CERTIFIED APPLIANCE**

Intel® Select Solution configurations and benchmark results are
**INTEL VERIFIED**

# THE INTEL DIFFERENTIATION

TRANSISTORS & PACKAGING

ARCHITECTURE

MEMORY

INTERCONNECTS

SECURITY

SOFTWARE & SOLUTIONS

# JIM KELLER

SR. VICE PRESIDENT
GM, SILICON ENGINEERING GROUP
INTEL

DATA-CENTRIC
INNOVATION SUMMIT

# SNEAK PEEK INTO THE FUTURE

| 2018 | 2019 | 2020 |
| --- | --- | --- |
| | 14NM/10NM PLATFORM | |
| **CASCADE LAKE**<br>14NM<br>SHIPPING Q4'18 | **COOPER LAKE**<br>14NM | **ICE LAKE**<br>10NM |
| INTEL OPTANE PERSISTENT MEMORY<br>INTEL DLBOOST: VNNI<br>SECURITY MITIGATIONS | NEXT GEN INTEL DLBOOST: BFLOAT16 | |

**LEADERSHIP PERFORMANCE**

# SUMMARY

## IT'S A NEW ERA OF DATA-CENTRIC COMPUTING

FUELED BY CLOUD, NETWORK | 5G | EDGE, ARTIFICIAL INTELLIGENCE

## THE DATA-CENTRIC OPPORTUNITY IS MASSIVE

LARGEST OPPORTUNITY IN INTEL'S HISTORY, OVER $200B TAM BY 2022

## INTEL HAS UNPARALLELED ASSETS TO FUEL GROWTH

PORTFOLIO OF LEADERSHIP PRODUCTS TO MOVE, STORE AND PROCESS DATA

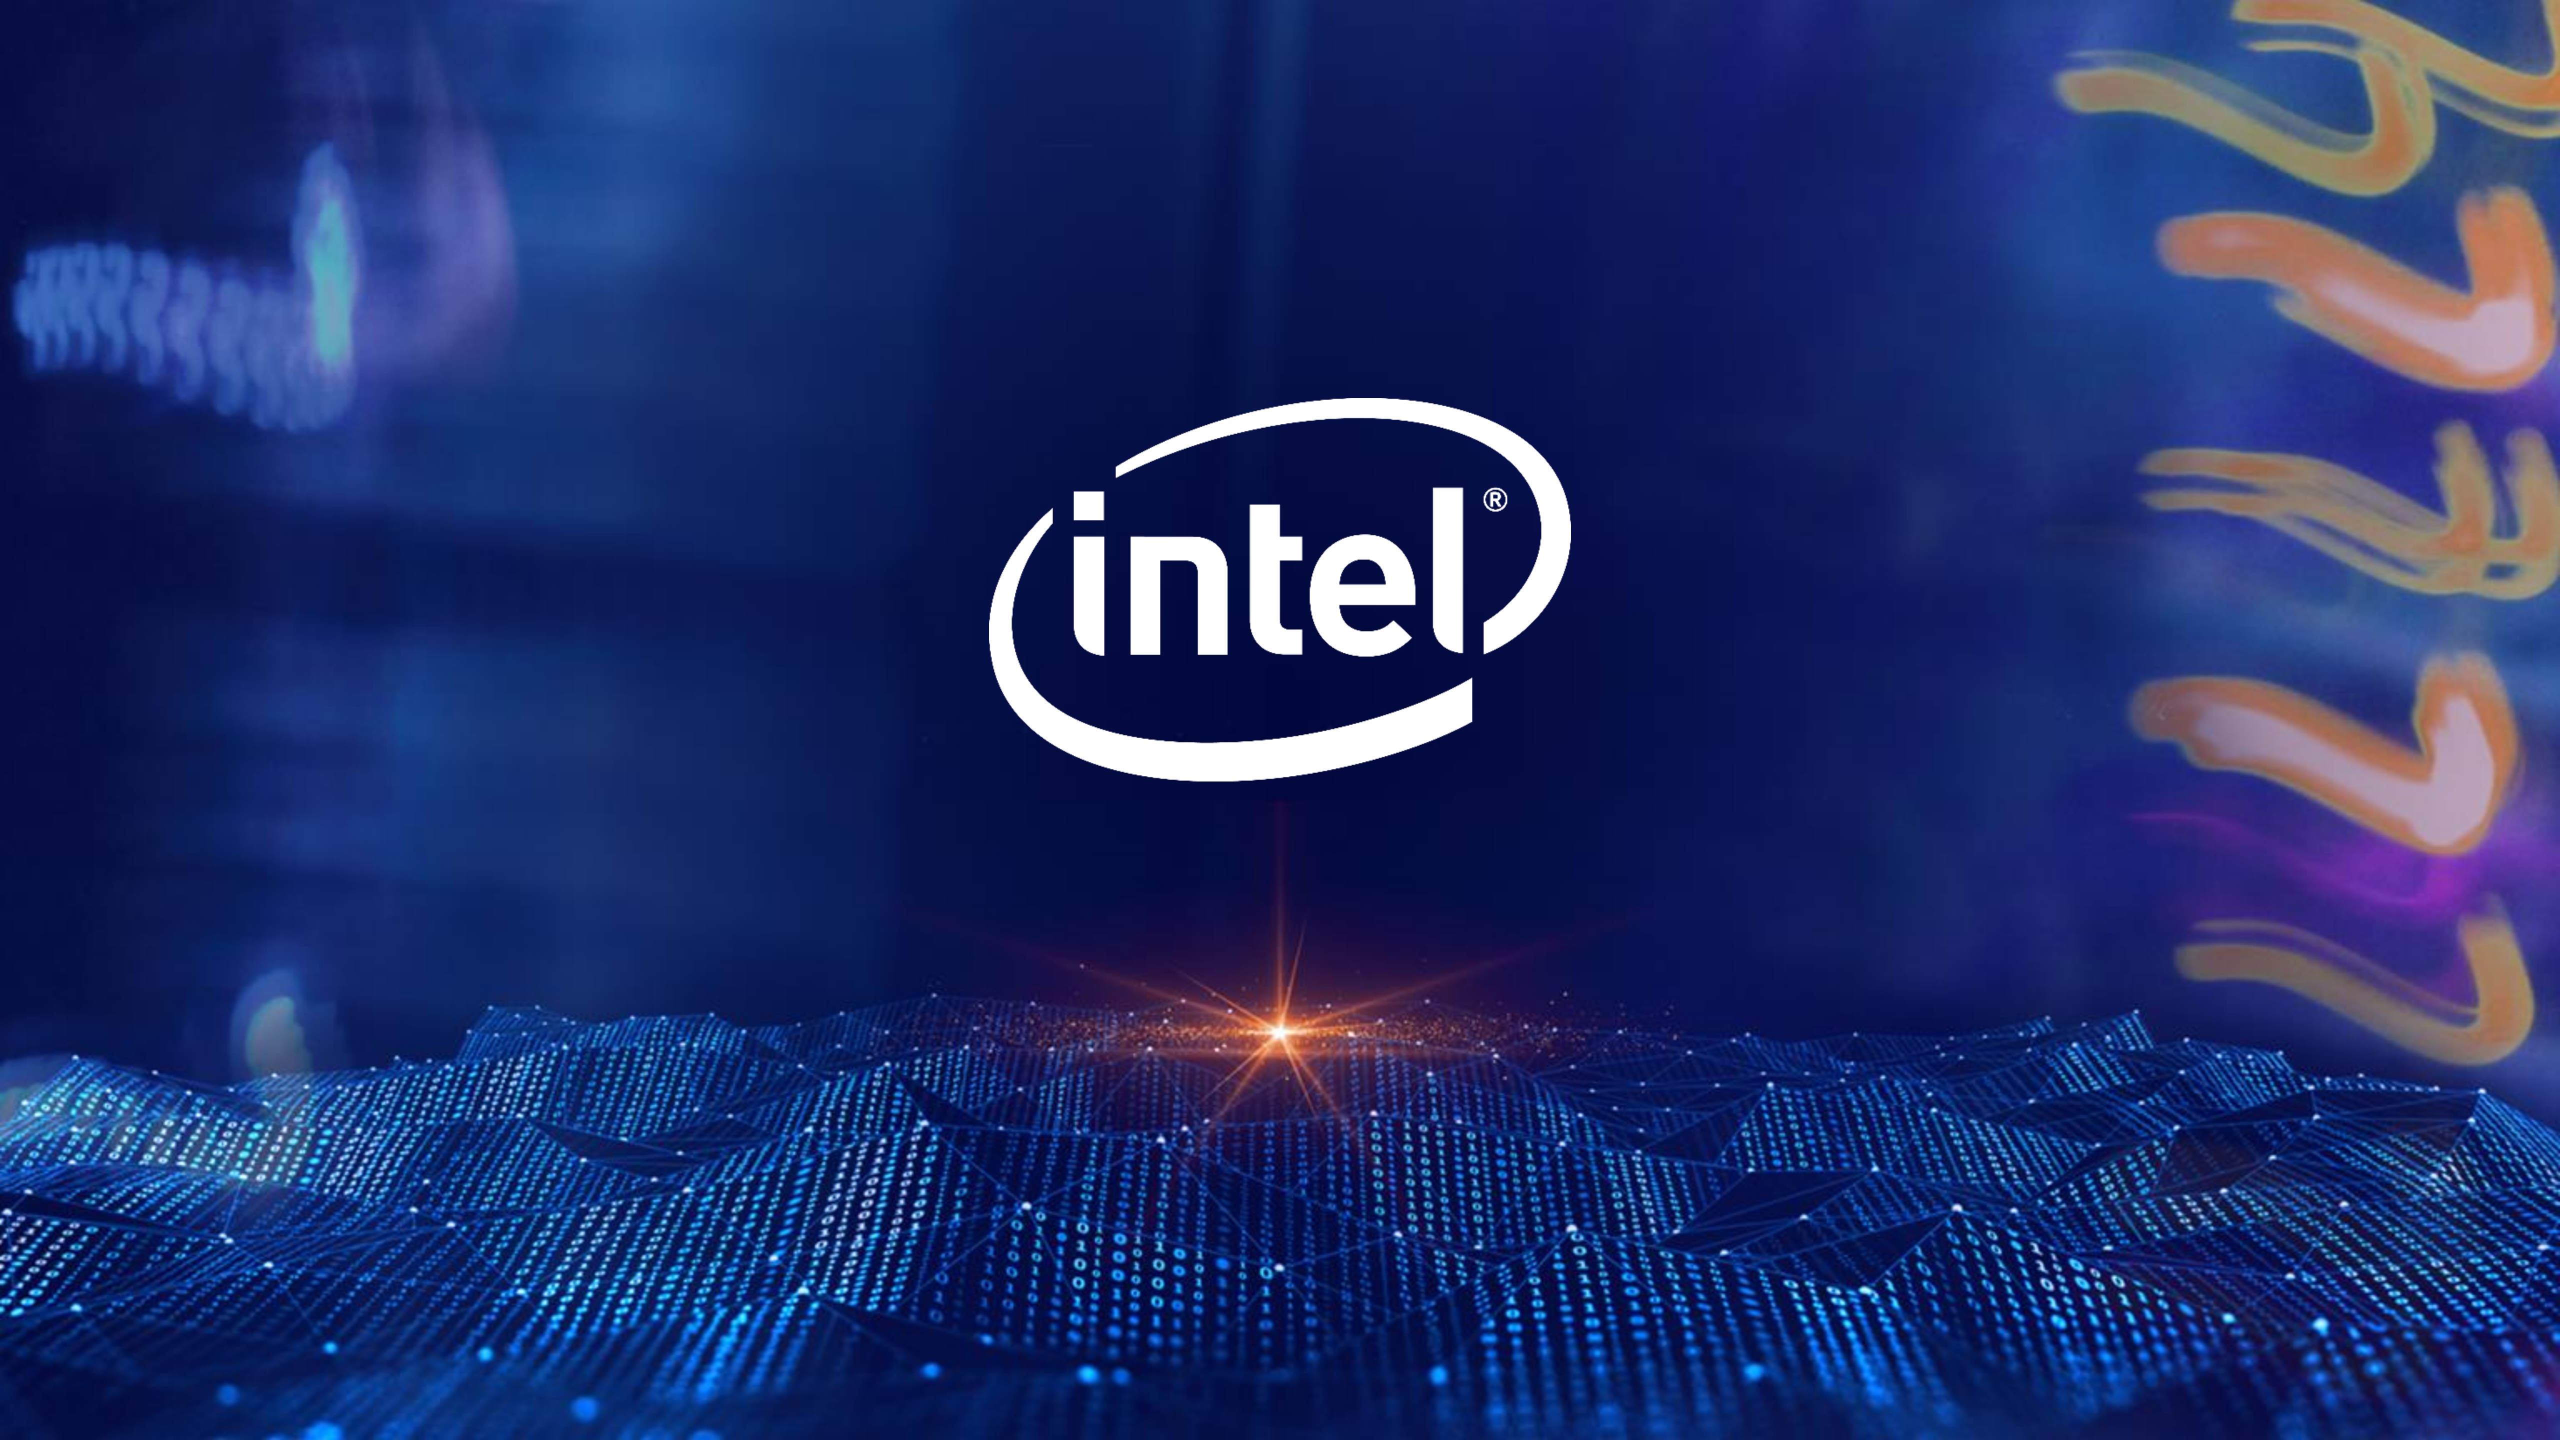
intel

# 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.

# CONFIGURATION DETAILS

**1.48x: Per Core Performance**

Intel Xeon Platinum 8180: Intel Xeon-based Reference Platform with 2 Intel Xeon 8180 (2.5GHz, 28 core) processors, BIOS ver SE5C620.86B.00.01.0014.070920180847, 07/09/2018, microcode: 0x200004d, HT ON, Turbo ON, 12x32GB DDR4-2666, 1 SSD, Ubuntu 18.04.1 LTS (4.17.0-041700-generic Retpoline), 1-copy SPEC CPU 2017 integer rate base benchmark compiled with Intel Compiler 18.0.2 -O3, executed on 1 core using taskset and numactl on core 0. Estimated score = 6.59, as of 8/2/2018 tested by Intel

AMD EPYC 7601: Supermicro AS-2023US-TR4 with 2S AMD EPYC 7601 with 2 AMD EPYC 7601 (2.2GHz, 32 core) processors, BIOS ver 1.1a, 4/26/2018, microcode: 0x8001227, SMT ON, Turbo ON, 16x32GB DDR4-2666, 1 SSD, Ubuntu 18.04.1 LTS (4.17.0-041700-generic Retpoline), 1-copy SPEC CPU 2017 integer rate base benchmark compiled with AOCC ver 1.0 -Ofast, -march=znver1, executed on 1 core using taskset and numactl on core 0. Estimated score = 4.45, as of 8/2/2018 tested by Intel

**3.20x: High Performance Linpack**

Intel Xeon Platinum 8180: Intel Xeon-based Reference Platform with 2 Intel Xeon 8180 (2.5GHz, 28 core) processors, BIOS ver SE5C620.86B.00.01.0014.070920180847, 07/09/2018, microcode: 0x200004d, HT ON (1 thread per core), Turbo ON, 12x32GB DDR4-2666, 1 SSD, Ubuntu 18.04.1 LTS (4.17.0-041700-generic Retpoline), High Performance Linpack v2.1, compiled with Intel(R) Parallel Studio XE 2018 for Linux, Intel MPI and MKL Version 18.0.0.128, Benchmark Config: Nb=384, N=203136, P=1, Q=2, Q=4, Score = 3507.38GFs, as of July 31, 2018 tested by Intel

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, as of July 31, 2018 tested by Intel

**1.85x: Database**

Intel Xeon Platinum 8180: Intel Xeon-based Reference Platform with 2 Intel Xeon 8180 (2.5GHz, 28 core) processors, BIOS ver SE5C620.86B.0X.01.0115.012820180604, microcode: 0x2000043, HT ON, Turbo ON, 24x32GB DDR4-2666, 1 x Intel DC P3700 PCI-E SSD (2TB, 1/2 Height PCIe 3.0, 20nm, MLC), Red Hat Enterprise Linux 7.4 (3.10.0-693.11.6.el7.x86_64 IBRS), HammerDB ver 2.3, PostgreSQL ver 9.6.5, Score = 2,250,481 tpm, as of 3/15/2018 tested by Intel

AMD EPYC 7601: HPE Proliant DL385 Gen10 with 2 AMD EPYC 7601 (2.2GHz, 32 core) processors, ROM ver 1.06, microcode: 0x8001227, SMT ON, Turbo ON, 16x32GB DDR4-2666, 1 x Intel DC P3700 PCI-E SSD (2TB, 1/2 Height PCIe 3.0, 20nm, MLC), Red Hat Enterprise Linux 7.4 (3.10.0-693.21.1.el7.x86_64 Retpoline), HammerDB ver 2.3, PostgreSQL ver 9.6.5, Score = 1,210,575 tpm, as of 4/12/2018 tested by Intel

**1.45x: Memcached (Memory Object Caching)**

Intel Xeon Platinum 8180: Intel Reference Platform with 2 Intel Xeon 8180 (2.5GHz, 28C) processors, BIOS ver SE5C620.86B.00.01.0014.070920180847, 07/09/2018, microcode: 0x200004d, HT ON, Turbo ON, 12x32GB DDR4-2666, 1SSD, 1 40GbE PCIe XL710 Adapter, Ubuntu 18.04.1 LTS (4.17.0-041700-generic Retpoline), Memcached using YCSB benchmark Workloadc, YCSB 0.16.0, Memcached v1.5.9, Max throughput (ops/sec) with P99 latency < 1ms, Score: 2711265 ops/sec, as of 8/2/2018 tested by Intel

AMD EPYC 7601: Supermicro AS-2023US-TR4 with 2 AMD EPYC 7601 (2.2GHz, 32C) processors, BIOS ver 1.1a, 4/26/2018, microcode: 0x8001227, SMT ON, Turbo ON, 16x32GB DDR4-2666, 1SSD, 1 40GbE PCIe XL710 Adapter, Ubuntu 18.04 LTS, (4.17.0-041700-generic Retpoline), Memcached using YCSB benchmark Workloadc, YCSB 0.16.0, Memcached v1.5.9, Max throughput (ops/sec) with P99 latency < 1ms, Score: 1862841 ops/sec, as of 8/2/2018 tested by Intel

**1.72x: L3 Packet Forwarding**

Intel Xeon Platinum 8180: Supermicro X11DPG-QT with 2 Intel Xeon-SP 8180 (2.5GHz, 28C) processors, BIOS ver 2.0b, microcode: 0x2000043, 12x32GB DDR4-2666, 1 SSD, 2x Intel XXV710-DA2 PCI Express (2x25GbE), DPDK L3fwd sample application (IPv4 LPM, 256B packet size, 625000 flows), DPDK 17.11, Ubuntu 17.10, (4.13.0-31-generic IBRS), HT ON, Turbo OFF, Score= 42.22 Million Packets / second, as of 8/2/2018 tested by Intel

AMD EPYC 7601, Supermicro AS-2023US-TR4 with 2 AMD EPYC 7601 (2.2GHz, 32C) processors, BIOS ver 1.1a, microcode: 0x8001227, 16x32GB DDR4-2666, 1 SSD, 2x Intel XXV710-DA2 PCI Express (2x25GbE), DPDK L3fwd sample application (IPv4 LPM, 256B packet size, 625000 flows), DPDK 17.11, Ubuntu 17.10 (4.13.0-36-generic Retpoline), SMT ON, Turbo (core boost) OFF, Score= 24.52 Million Packets / second, as of 8/2/2018 tested by Intel

# INTEL OPTANE PERSISTENT MEMORY CONFIGURATION DETAILS

Performance results are based on testing: 8X (8/2/2018), 9X Reads/11X Users (5/24/2018), Minutes to Seconds (5/30/2018) and may not reflect all publicly available security updates. No product can be absolutely secure.
Results have been estimated based on tests conducted on pre-production systems: 8x (running OAP with 2.6TB scale factor on IO intensive queries), 9X Reads/11X Users (running Cassandra optimized for persistent memory), and Minutes to Seconds (running Aerospike* Hybrid Memory Architecture optimized for persistent memory), and provided to you for informational purposes.

# AI PERFORMANCE CONFIGURATION DETAILS

**1x inference throughput improvement in 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, dummy 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“.

**2.8x inference throughput improvement in January 2018:**

Tested by Intel as of Jan 19th 2018 Processor :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, SSD sda RS3WC080 HDD 744.1GB,sdb RS3WC080 HDD 1.5TB,sdc RS3WC080 HDD 5.5TB , Deep Learning Framework Intel® Optimization for caffe version:f6d01efbe93f70726ea3796a4b89c612365a6341 Topology::resnet_50_v1 BIOS:SE5C620.86B.00.01.0009.101920170742 MKLDNN: version: ae00102be506ed0fe2099c6557df2aa88ad57ec1 NoDataLayer. . Datatype:FP32 Batchsize=64 Measured: 652.68 imgs/sec vs 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, dummy 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“.

**5.4x inference throughput improvement in August 2018:**

Tested by Intel as of measured July 26th 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:a3d5b022fe026e9092fc7abc7654b1162ab9940d Topology::resnet_50_v1 BIOS:SE5C620.86B.00.01.0013.030920180427 MKLDNN: version:464c268e544bae26f9b85a2acb9122c766a4c396 instances: 2 instances : https://software.intel.com/en-us/articles/boosting-deep-learning-training-inference-performance-on-xeon-and-xeon-phi NoDatasocket:2 (Results on Intel® Xeon® Scalable Processor were measured running multiple instances of the framework. Methodology described hereLayer. Datatype: INT8 Batchsize=64 Measured: 1233.39 imgs/sec 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, dummy 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“.

**11X inference thoughput improvement with CascadeLake**:

Future Intel Xeon Scalable processor (codename Cascade Lake) results have been estimated or simulated using internal Intel analysis or architecture simulation or modeling, and provided to you for informational purposes. Any differences in your system hardware, software or configuration may affect your actual performance 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, dummy 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“.