TRANSFORMING STORAGE WITH INNOVATIONS IN NON-VOLATILE MEMORY

Bill Bollengier, Director

Intel Non-volatile Memory Solutions Group

WE ARE IN A DATA-CENTRIC WORLD

All data must be

stored, processed, analyzed

?????

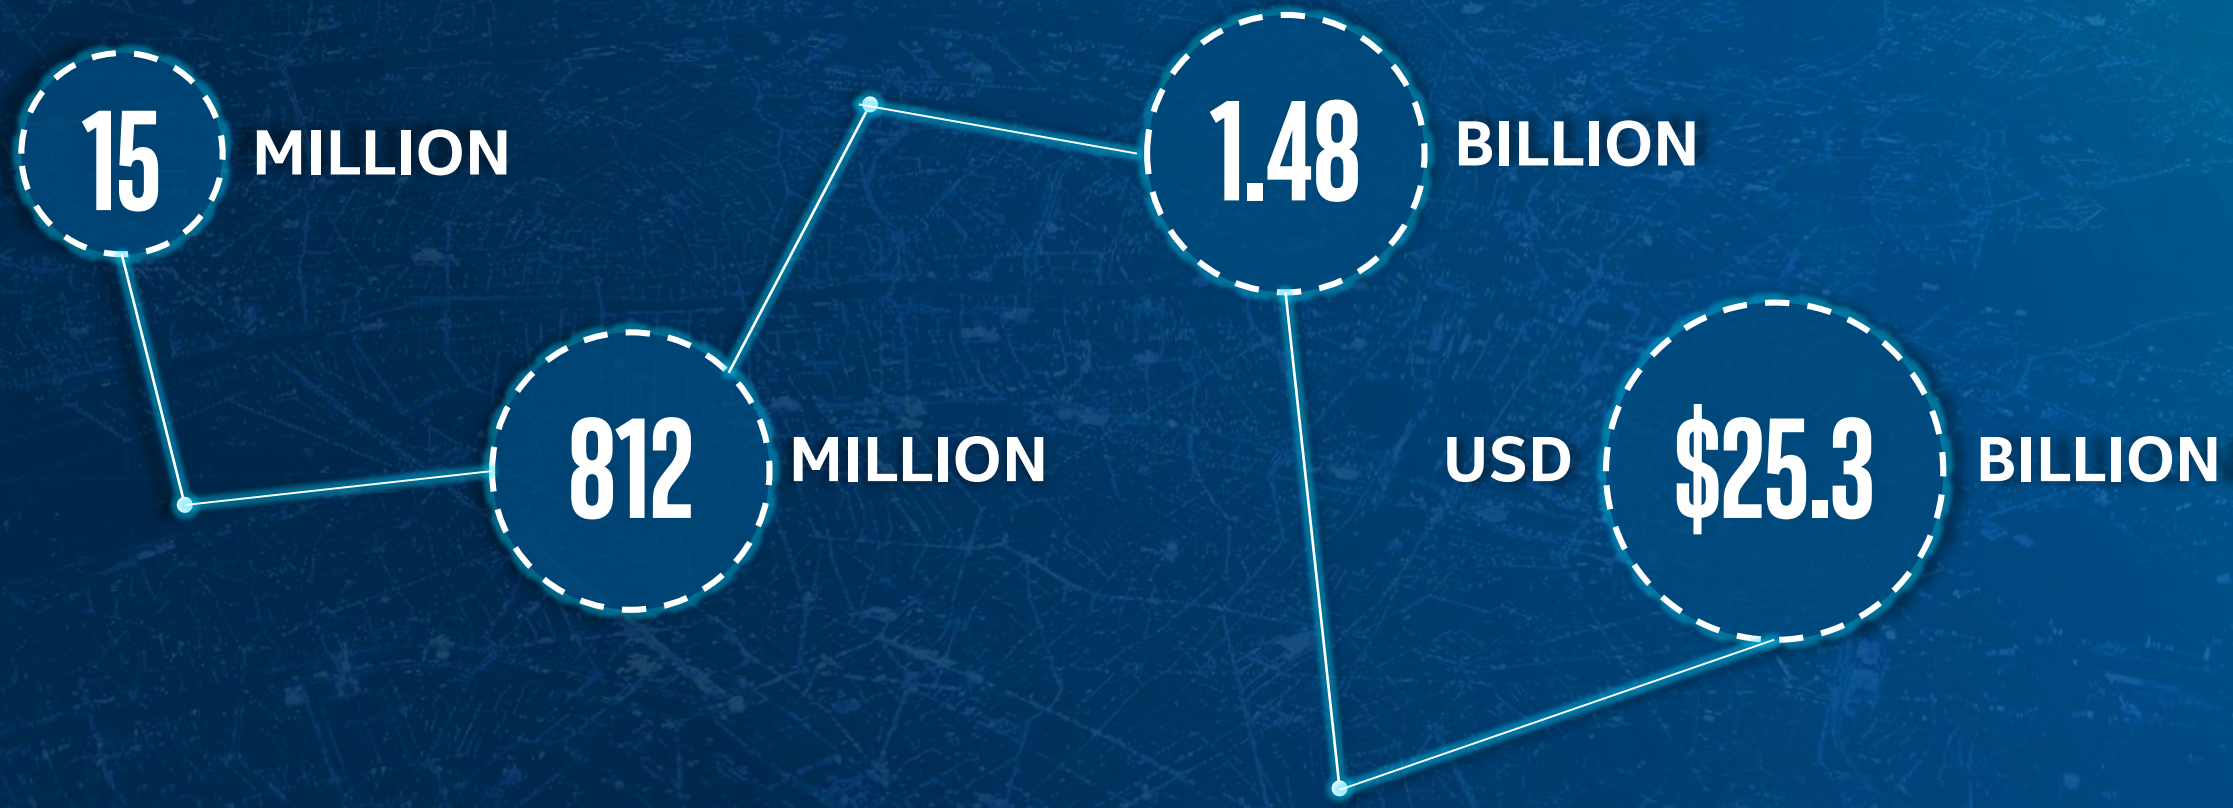
15
MILLION
812
MILLION
1.48
BILLION
USD
$25.3
BILLION

ALIBABA* 'SINGLES DAY' 2017

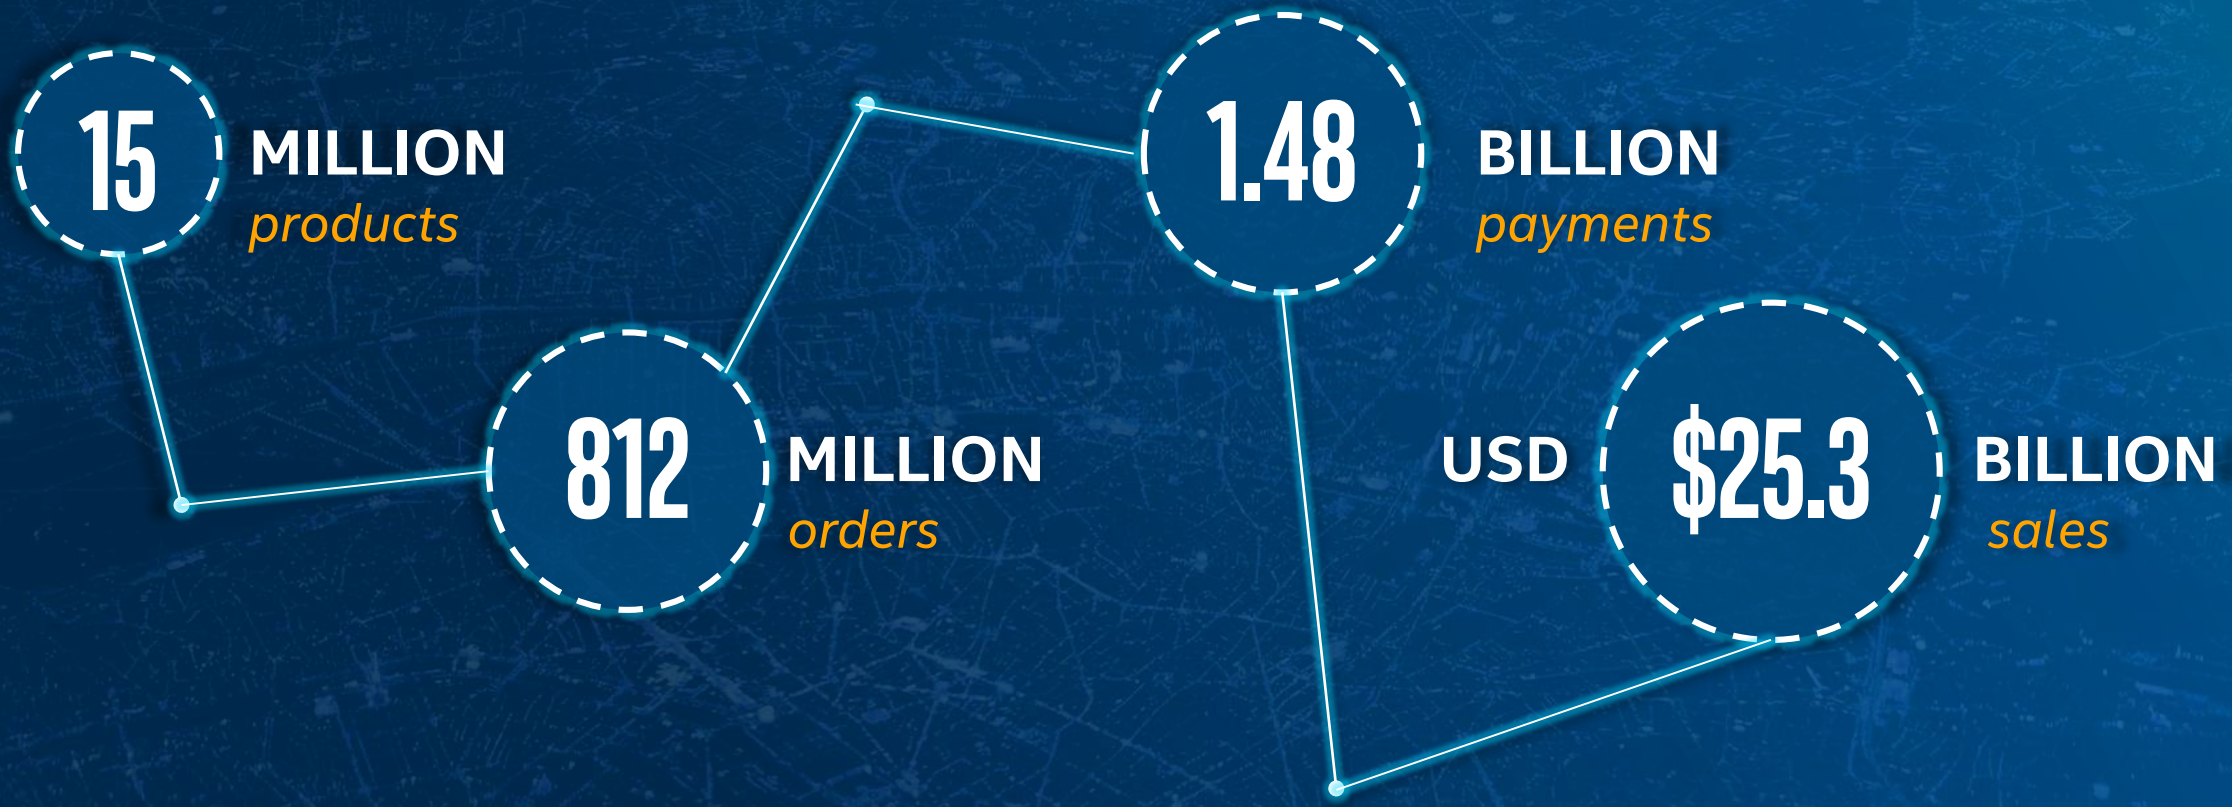

*Other names and brands may be claimed as the property of others
Source: www.forbes.com/sites/helenwang/2017/11/12/alibabas-singles-day-by-the-numbers-a-record-25-billion-haul/#52d3a77d1db1

DATA GROWTH IS HAPPENING ALL OVER

10 GB[1]
AVERAGE INTERNET USER

3,000 GB[2]
SMART HOSPITAL

4,000 GB[3]
AUTONOMOUS DRIVING

40,000 GB[2]
AIRPLANE DATA

1 M GB[2]
SMART FACTORY

1.Source: http://www.cisco.com/c/en/us/solutions/service-provider/vni-network-traffic-forecast/infographic.html
2.Source: http://www.cisco.com/c/en/us/solutions/collateral/service-provider/global-cloud-index-gci/Cloud_Index_White_Paper.html
3.Source: https://dataflog.com/read/self-driving-cars-create-2-petabytes-data-annually/172

VIRTUOUS CYCLE

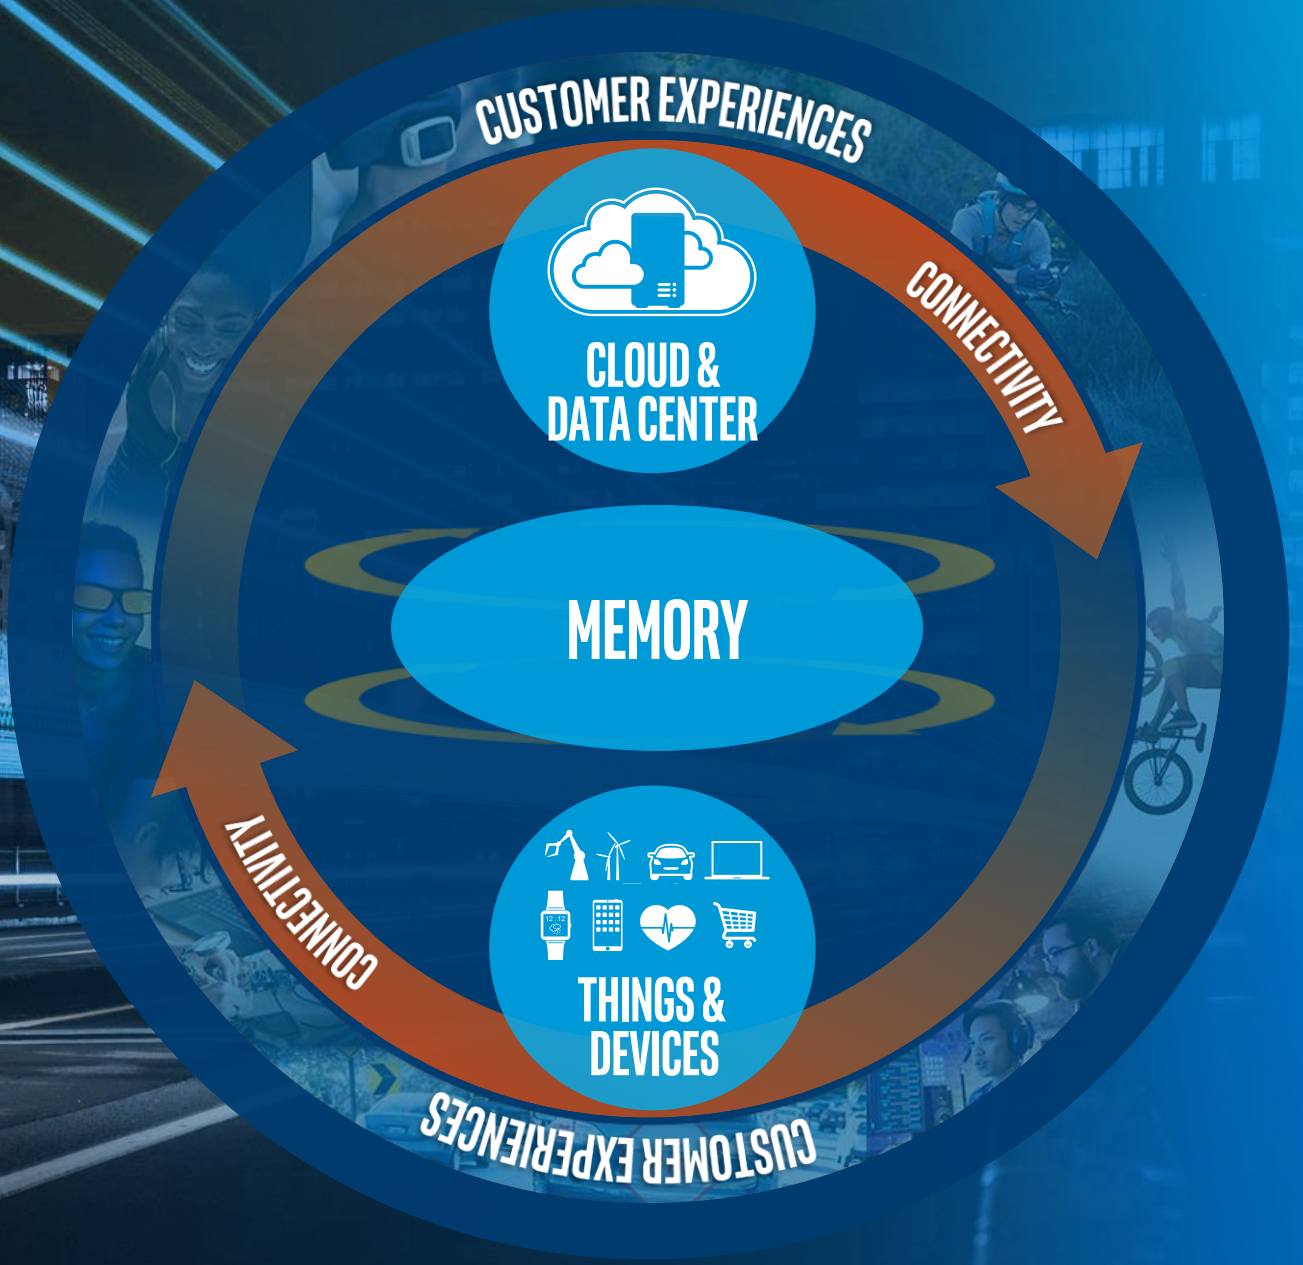

DATA IS STORED BY DIFFERENT TIERS

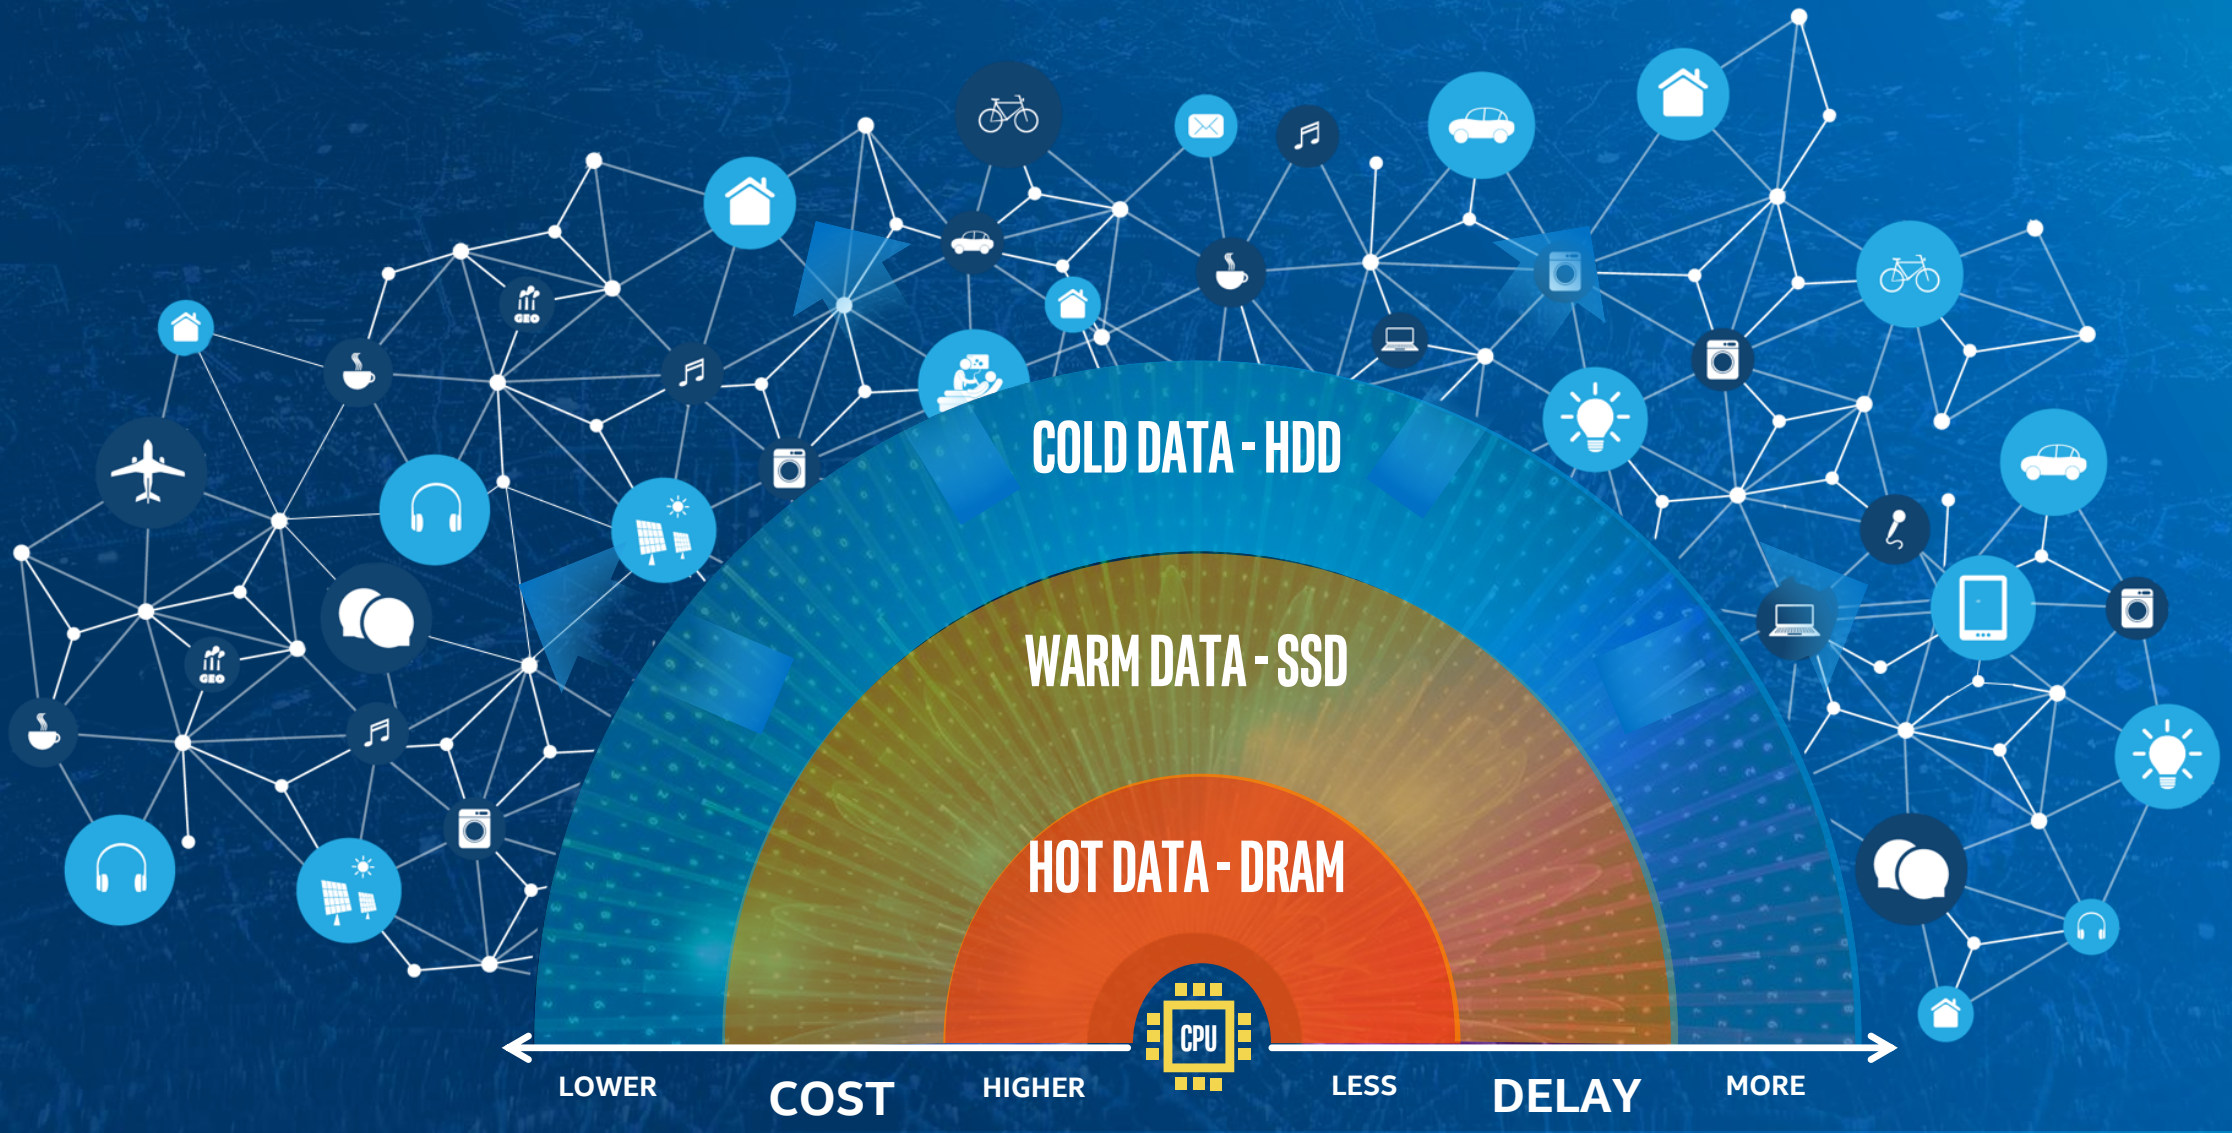

INTEL IS INVESTING IN 2 TECHNOLOGIES

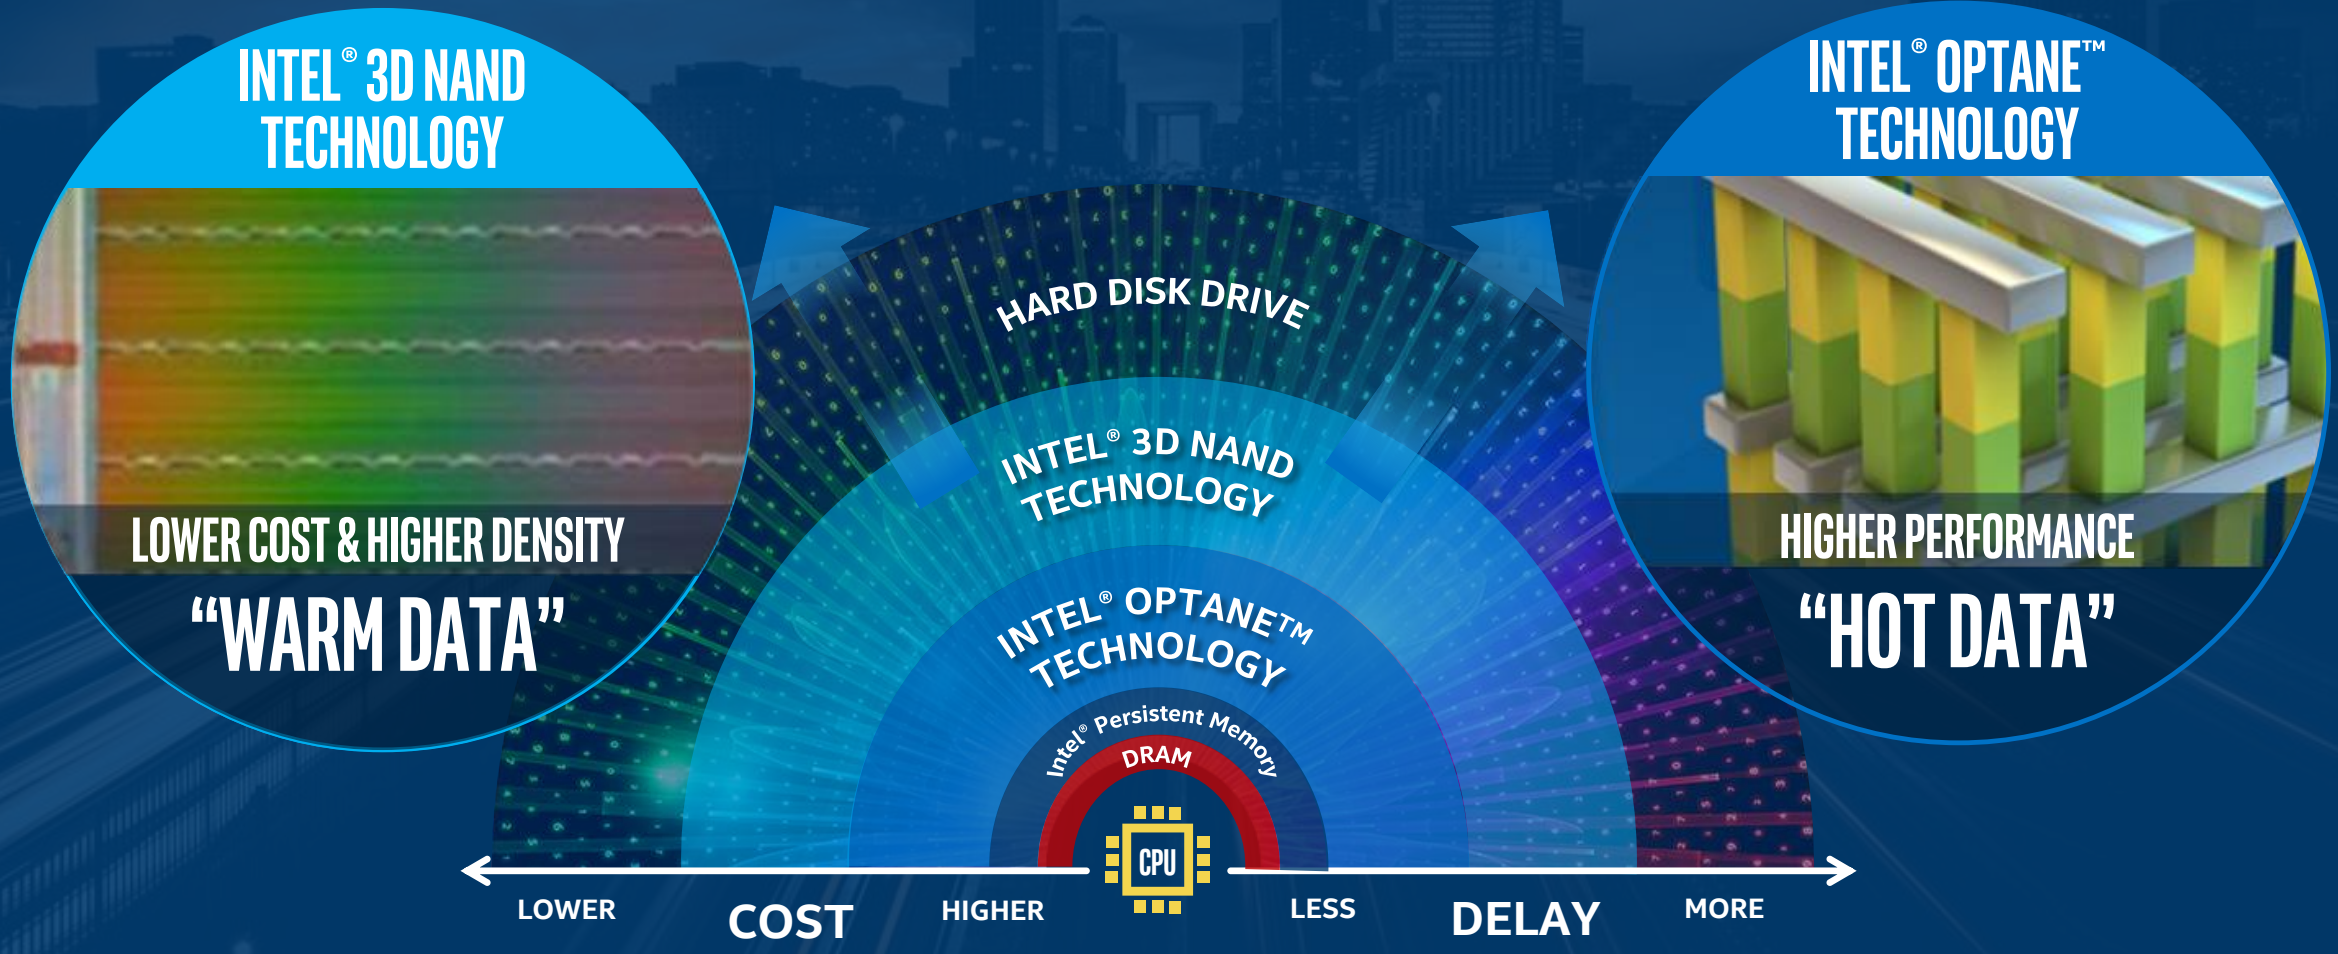

INTEL® OPTANE™ TECHNOLOGY

Highest Performance: Break The Bottleneck

QOS

THROUGHPUT (IOPS)

LATENCY

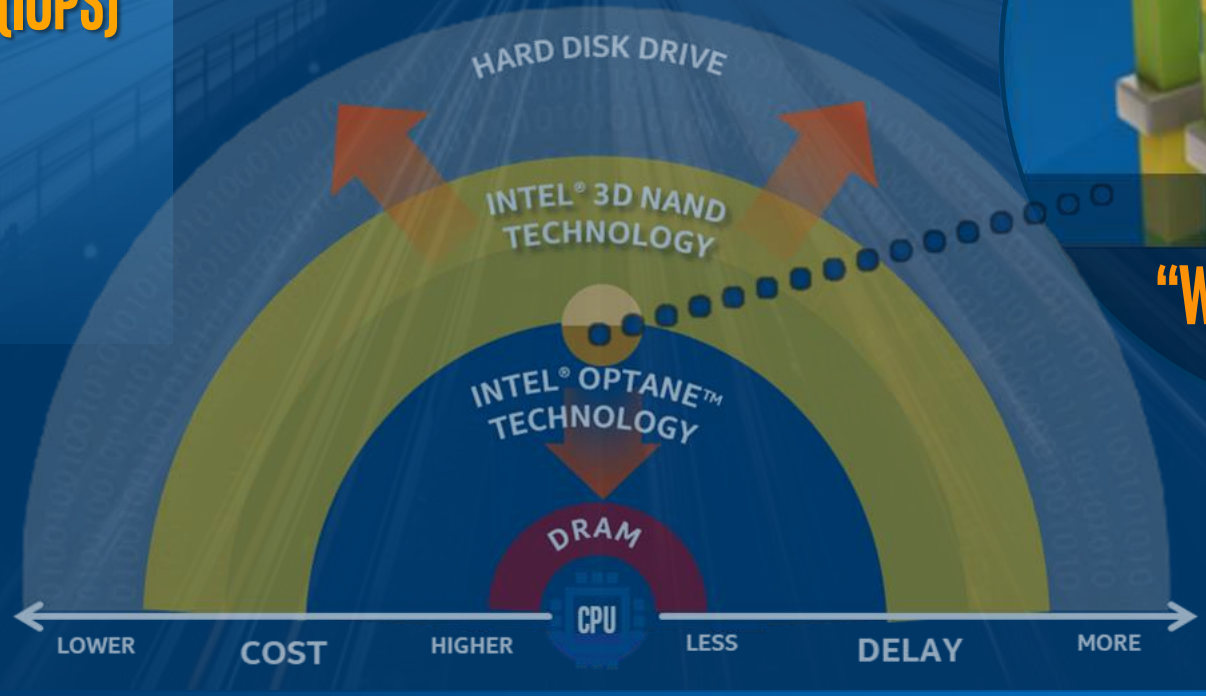

INTEL® OPTANE™ TECHNOLOGY

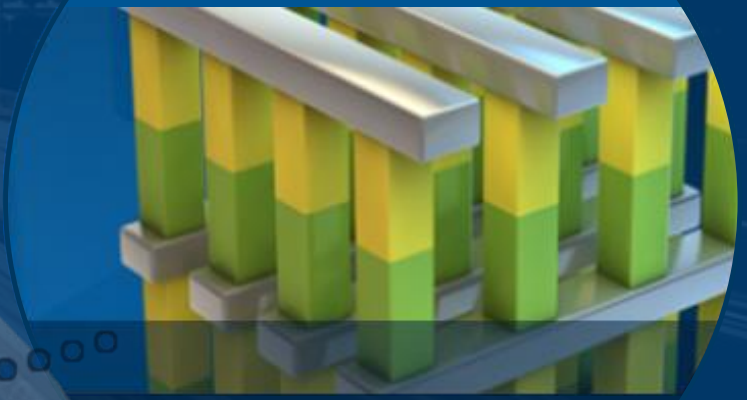

"WORKING STORAGE"

HIGHER PERFORMANCE

THE VALUE IS IN THE GAP

MEMORY

SRAM
Latency: 1X
Size of Data: 1X

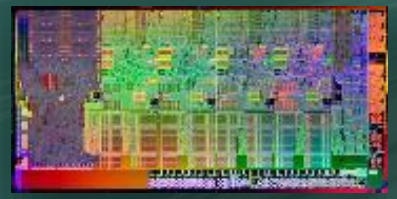

DRAM
Latency: ~10X
Size of Data: ~100X

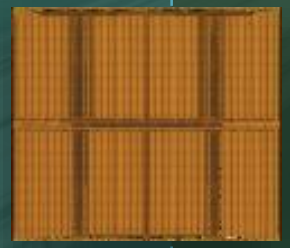

New experiences

New levels of scale

Accelerate existing apps

STORAGE

NAND SSD
Latency: ~100,000X
Size of Data: ~1,000X

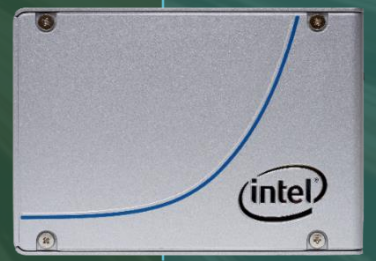

HDD
Latency: ~10 Million X
Size of Data: ~10,000X

Technology claims are based on comparisons of latency, density and write cycling metrics amongst memory technologies recorded on published specifications of in-market memory products against internal Intel specifications.

INTEL® 3D NAND TECHNOLOGY

Highest Density & Lowest Cost

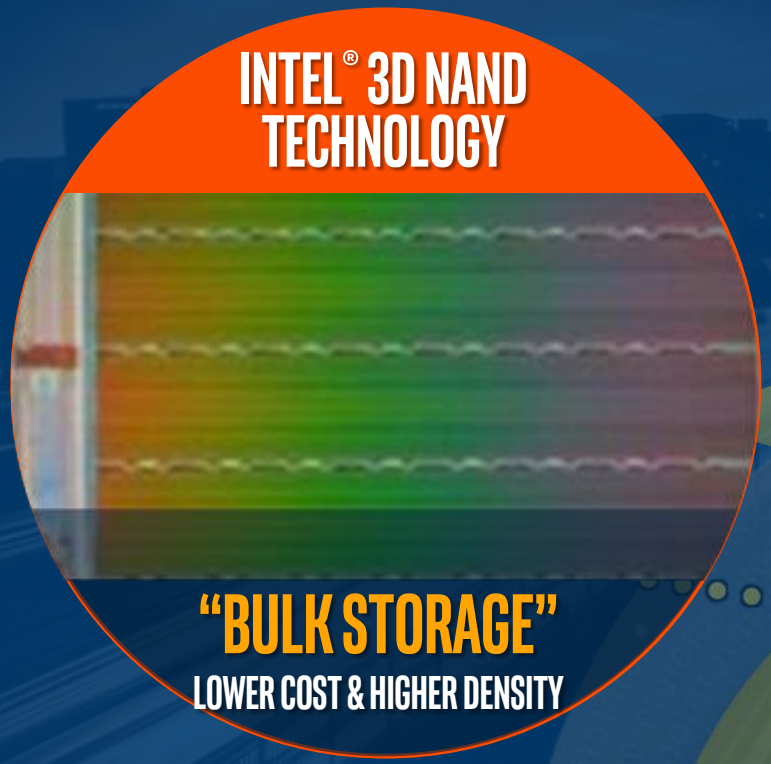

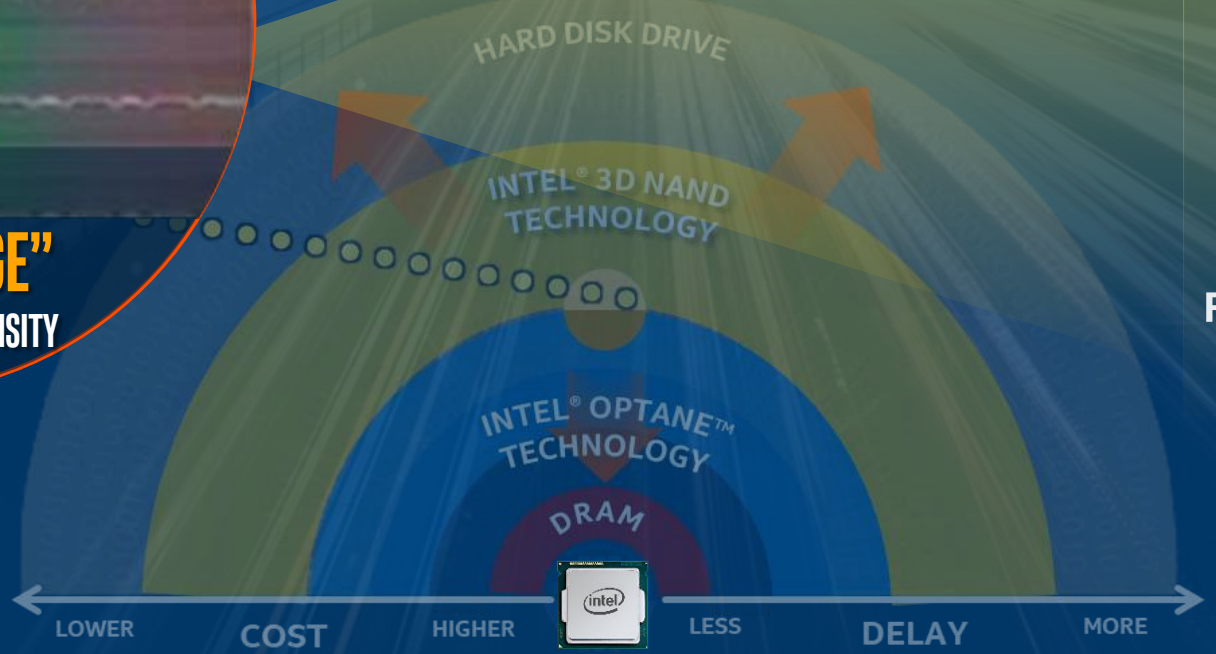

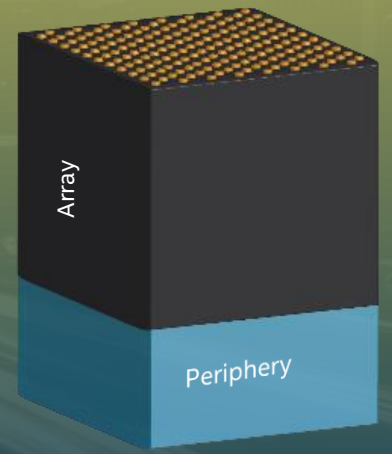

Intel® 3D NAND Floating Gate Technology

A PORTFOLIO OF SOLUTION COMPONENTS

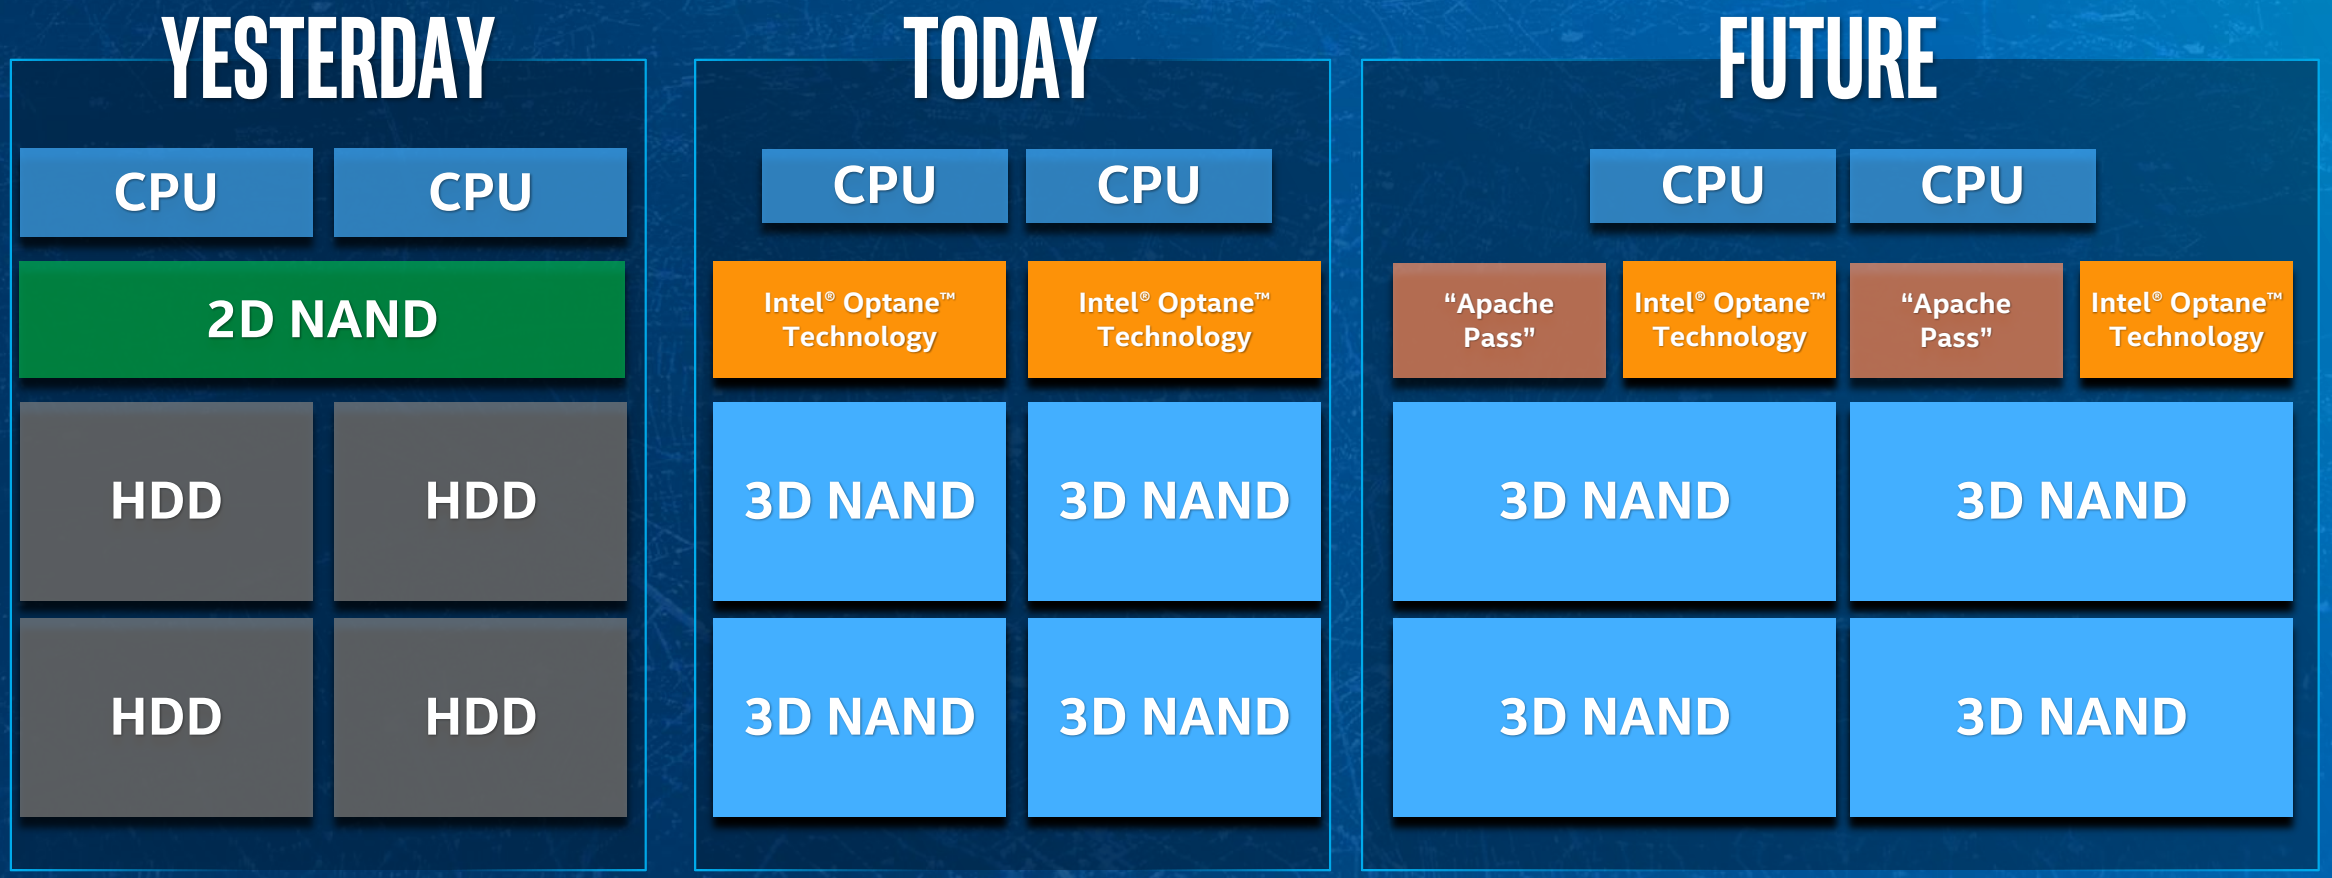

INTEL® OPTANE™ TECHNOLOGY FOR DATA CENTER CONFIGURATIONS

FAST STORAGE AND CACHING

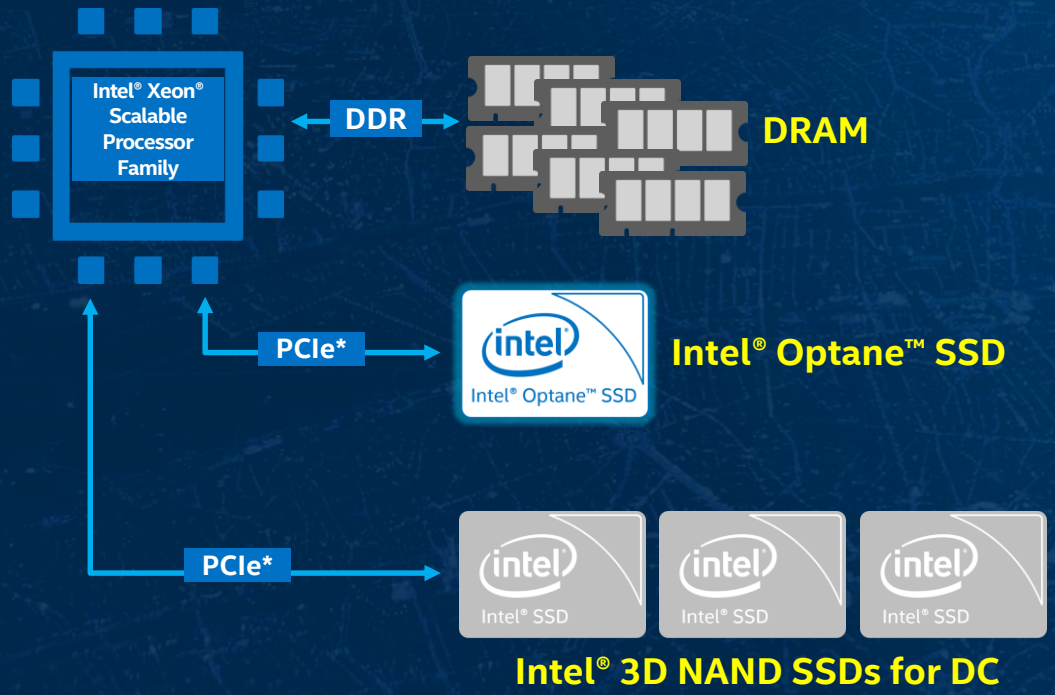

EXTEND MEMORY

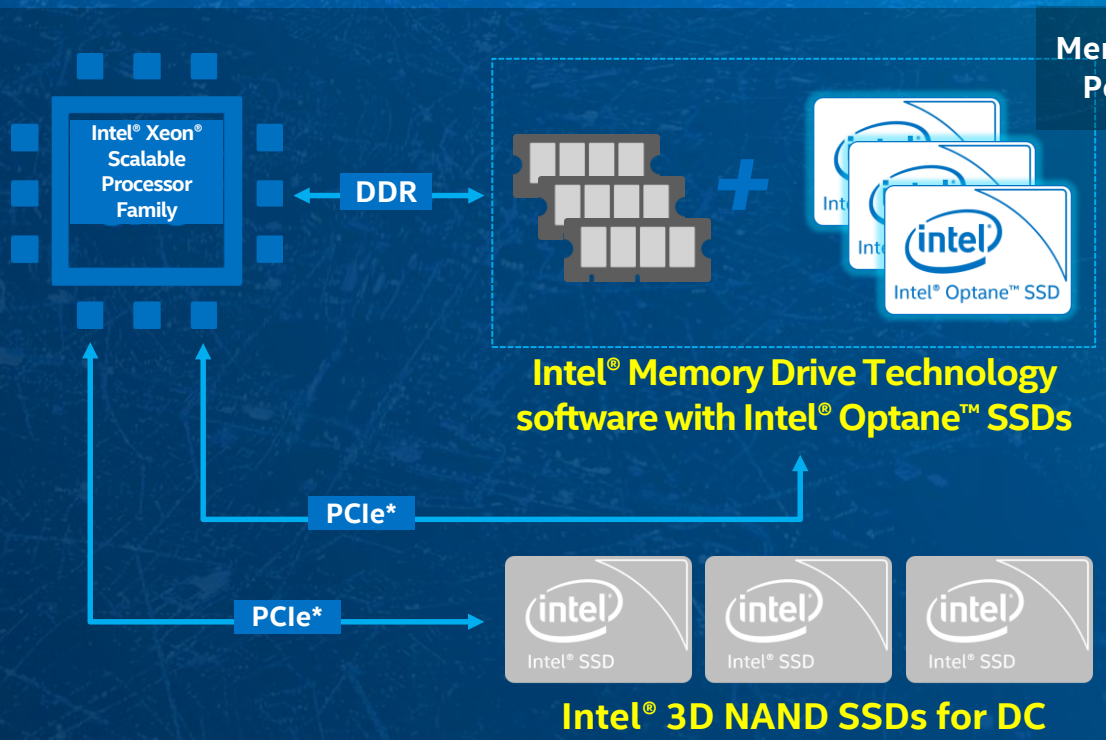

*Other names and brands may be claimed as the property of others.

ENABLING NEW USAGES & FORM FACTORS

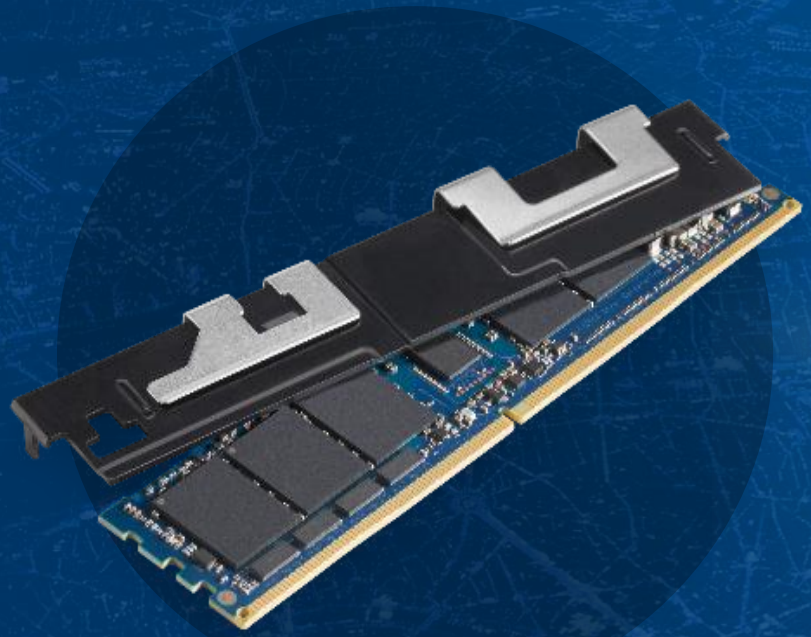

THE BEST OF BOTH WORLDS

DRAM ATTRIBUTES

Performance comparable to DRAM at *low latencies*[1]

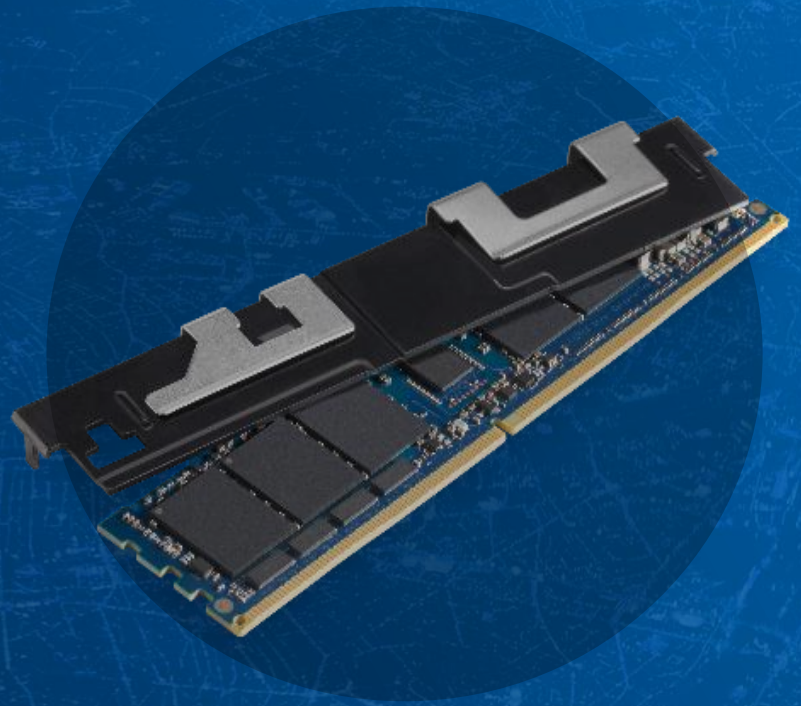

NAND SSD ATTRIBUTES

Data persistence with higher capacity than DRAM[2]

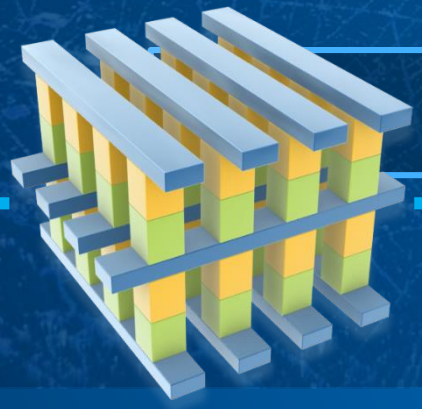

3D XPoint™ Technology

INTEL® OPTANE™ TECHNOLOGY + 3D NAND TECHNOLOGY

Plenty and affordable memory

High performance storage
(latency, bandwidth, QoS, endurance)

Application managed memory

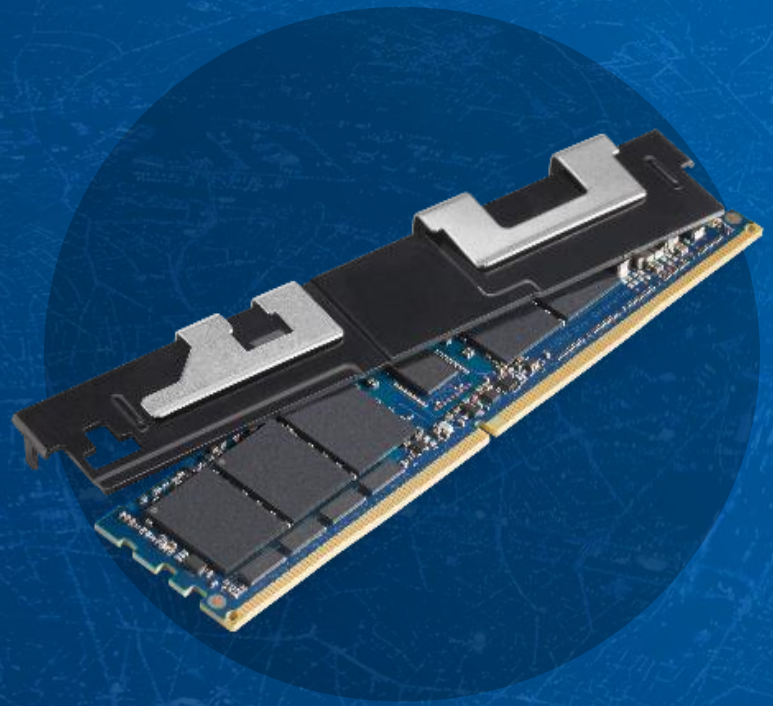

More and extended VMs

Capacity for In-Memory Database

Super-fast storage

Larger memory pools

PROVIDING LOWER AND CONSISTENT LATENCY WITH MORE CAPACITY PER DOLLAR

"RULER" FORM FACTOR

EDSFF Long

EDSFF Short

https://edsffspec.org

BUILT IN SERVICEABILITY

Programmable LEDs to quickly locate failed drives, offline drives, and unpopulated slots

Carrier-less design with integrated latch removes need for drive carriers

Enclosure Management with slot level **power control** enables single drive isolation or system level power loss

1PB IN 1U

1PB IN 42U
w/2 TB HDDs

STORAGE CAPACITY

1PB IN 1U
INTEL® 3D NAND SSD, 32TB RULER IN 2018

Opening up new use cases in warm storage with disruptive total cost of ownership

RESOURCES

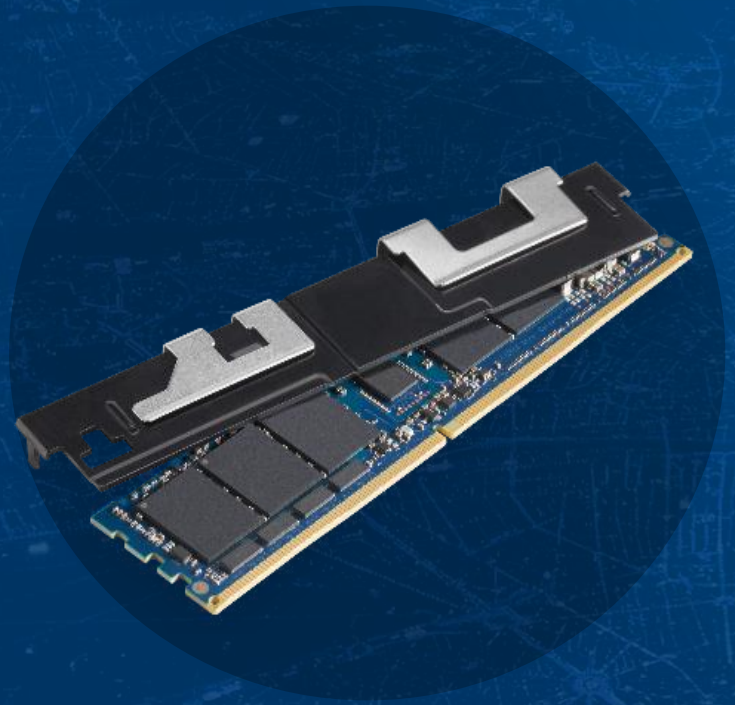

- pmem.io
- software.intel.com/en-us/persistent-memory

- edsffspec.org


THANK YOU!

LEGAL NOTICES AND DISCLAIMERS

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at intel.com, or from the OEM or retailer.

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. Consult other sources of information to evaluate performance as you consider your purchase.

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest forecast, schedule, specifications and roadmaps.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

Intel, the Intel logo, Optane, and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries.

APPENDIX

1.Source: http://www.cisco.com/c/en/us/solutions/service-provider/vni-network-traffic-forecast/infographic.html

2.Source: http://www.cisco.com/c/en/us/solutions/collateral/service-provider/global-cloud-index-gci/Cloud_Index_White_Paper.html

3.Source: https://datafloq.com/read/self-driving-cars-create-2-petabytes-data-annually/172

4.Source – Intel. Measured data. Comparing Intel 32 tier TLC NAND die (Intel® Pro6000p/600p Series) to Samsung 32 tier TLC NAND die.

5.Test and System Configuration: Intel® Server Board S2600WTTR, Intel® Xeon® E5-2699 v3, Speed: 2.30GHz, Intel BIOS: Internal Release, DRAM: DDR3 – 32GB, OS: Linux* Centos* 7.0 kernel 4.8, Intel® SSD DC P4500 2TB. Testing performed by Intel. Comparing with Intel® SSD DC S3520 1.6TB, http://www.intel.com/content/www/us/en/solid-state-drives/solid-state-drives-dc-s3520-series.html. Testing performed by Intel.

6.10X Lower AFR: Source - Intel. Intel SSD Annualized Fail Rate Report for all of 2015.

For footnote 7-11 below - Common configuration. Baseline (HDD) - Intel® Core™ i5-7500 Processor, 65W TDP, 4C4T, Turbo up to 3.8GHz, Memory: 2x4GB DDR4-2400, Storage: Western Digital* 1TB 7200RPM WD1003FZEX, Intel HD Graphics 630, OS: Windows* 10. Intel® Optane™ Memory. Same configuration as above with 16GB Intel® Optane™ Memory Module (Engineering Sample). Gaming workloads were tested with same configuration, except using a discrete graphics card (NVIDIA* (EVGA) GTX 1080) with and without 16GB Intel® Optane™ Memory Module. Tested with 16GB Intel® Optane™ Memory Engineering Samples. Results may vary in final product, but we have a high confidence level that there will be no significant differences in performance.

7.Everyday Tasks - SYSmark* 2014 SE - benchmark from the BAPCo* consortium that measures the performance of Windows* platforms. SYSmark* tests four usage scenarios: Office Productivity, Media Creation, Data/Financial Analysis, and Responsiveness. SYSmark* contains real applications from Independent Software Vendors such as Microsoft* and Adobe*.

8.Browser Launch Workload – Workload developed by Intel® measuring the time elapsed to launch Google* Chrome

9.Game Launch & Level Load Workload – Workload developed by Intel® measuring the time elapsed to launch Bethesda Softworks* Fallout 4 and reach the Main Menu with intro videos disabled (Launch), and the time elapsed from the Main Menu to completion of level loading (Level Load)

10.Email Launch Workload – Workload developed by Intel® measuring the time elapsed to launch Microsoft* Outlook 2016 and load with a 250mb local data file

11.Windows* File Search – Workload developed by Intel® measuring the time elapsed using Microsoft* Windows File Search to locate a specified file in a non-indexed directory

12.Responsiveness defined as average read latency measured at queue depth 1 during 4k random write workload. Measured using FIO 2.15. Common configuration - Intel 2U PCSD Server ("Wildcat Pass"), OS CentOS 7.2, kernel 3.10.0-327.el7.x86_64, CPU 2 x Intel® Xeon® E5-2699 v4 @ 2.20GHz (22 cores), RAM 396GB DDR @ 2133MHz. Intel drives evaluated - Intel® Optane™ SSD DC P4800X 375GB and Intel® SSD DC P3700 1600GB. Samsung drives evaluated – Samsung® SSD PM1725a, Samsung® SSD PM1725, Samsung® PM963, Samsung® PM953. Micron drive evaluated – Micron® 9100 PCIe® NVMe™ SSD. Toshiba drives evaluated – Toshiba® ZD6300. Test – QD1 Random Read 4K latency, QD1 Random RW 4K 70% Read latency, QD1 Random Write 4K latency using fio-2.15.

13.System configuration: Common – 2 x 5-node Ceph clusters both on Ceph BlueStore Kraken release 11.0.2, each node with Ubuntu 16.04 updated to Linux kernel 4.6, each cluster using 4 x Intel® SSD DC P3520 2TB as OSD (object storage device). All NAND cluster with each node – Intel® SSD DC P3700 1.6TB for metadata (db + WAL), 2 x Intel® Xeon E5, NIC 4x10GbE. Cluster using Intel® Optane SSDs: each node – Intel® Optane™ SSD 187GB for metadata (db + WAL), 2 x Intel® Xeon® E5 Haswell, NIC 2x10GbE. Test – latency based on 16K RW from 100 clients with QD2 at >99.9% latency, performance based on 2 cluster RW (4K/8K/16K) results, NAND cluster limited by P3700, Optane scaling performance is estimation based on 4K RW data of Optane and P3520 SSD.

14.Xeon E5v4 All-DRAM memory configuration hardware limited up to 3TB (assumes 24 DIMM x 128GB). In a 2-socket CPU configuration, Intel® Memory Drive Technology software supports up to 20 x 375TB Intel® Optane™ SSD DC P4800X for a total addressable space of 6TB, while DRAM as a cache is only 3TB. Attainable capacity depends on server configuration. Please consult your server manufacturer.

15.Xeon E5v4 All-DRAM memory configuration hardware limited up to 3TB (assumes 24 DIMM x 128GB). In a 2-socket CPU configuration, Intel® Memory Drive Technology software supports up to 20 x 1.5TB Intel® Optane™ SSD DC P4800X for a total addressable space of 24TB, while DRAM as a cache is only 3TB. Attainable capacity depends on server configuration. Please consult your server manufacturer.


intel
experience
what's inside™