H　P　C

Huge　Paradigm　Change

２０１９年１２月１２日

インテル株式会社
ＨＰＣ事業開発部長　矢澤 克巳
ＨＰＣ テクニカル・ソリューション・セールス　力 翠湖

1ST ERA IN HPC

VERTICALLY INTEGRATED SYSTEMS

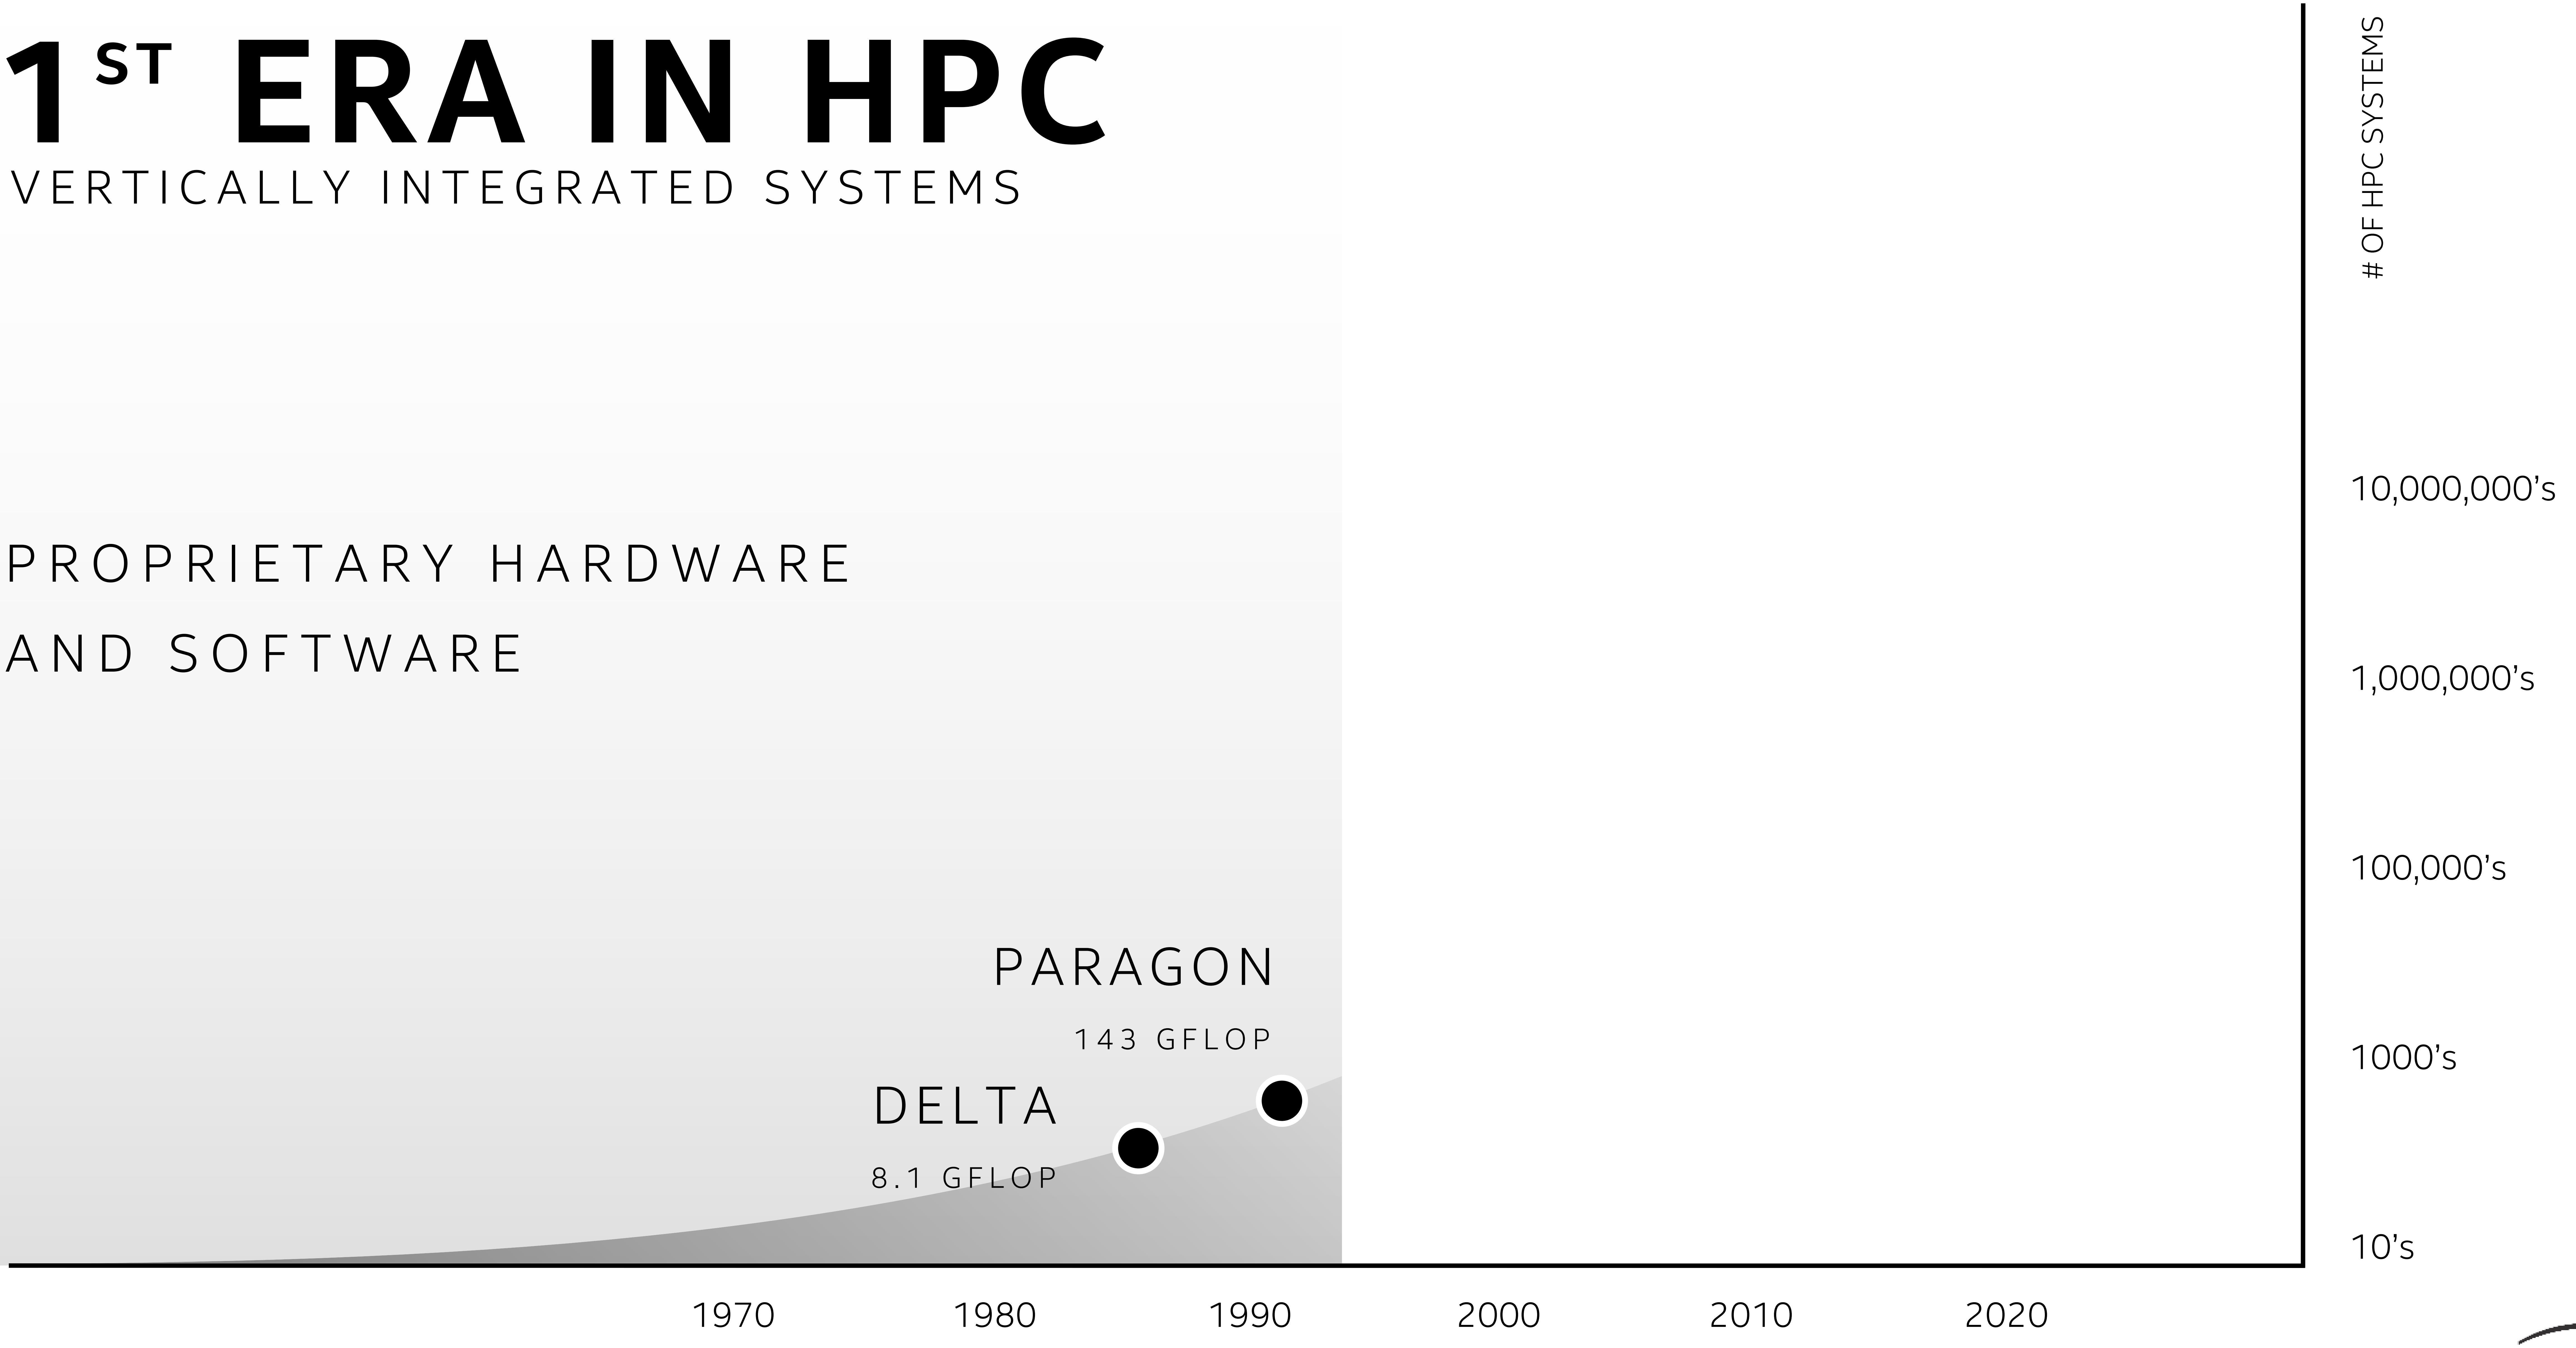

2ND ERA IN HPC

MOSTLY BASED ON GENERAL PURPOSE CPU'S

X86

LINUX

OPEN STANDARDS

BUILD TO ORDER SYSTEMS

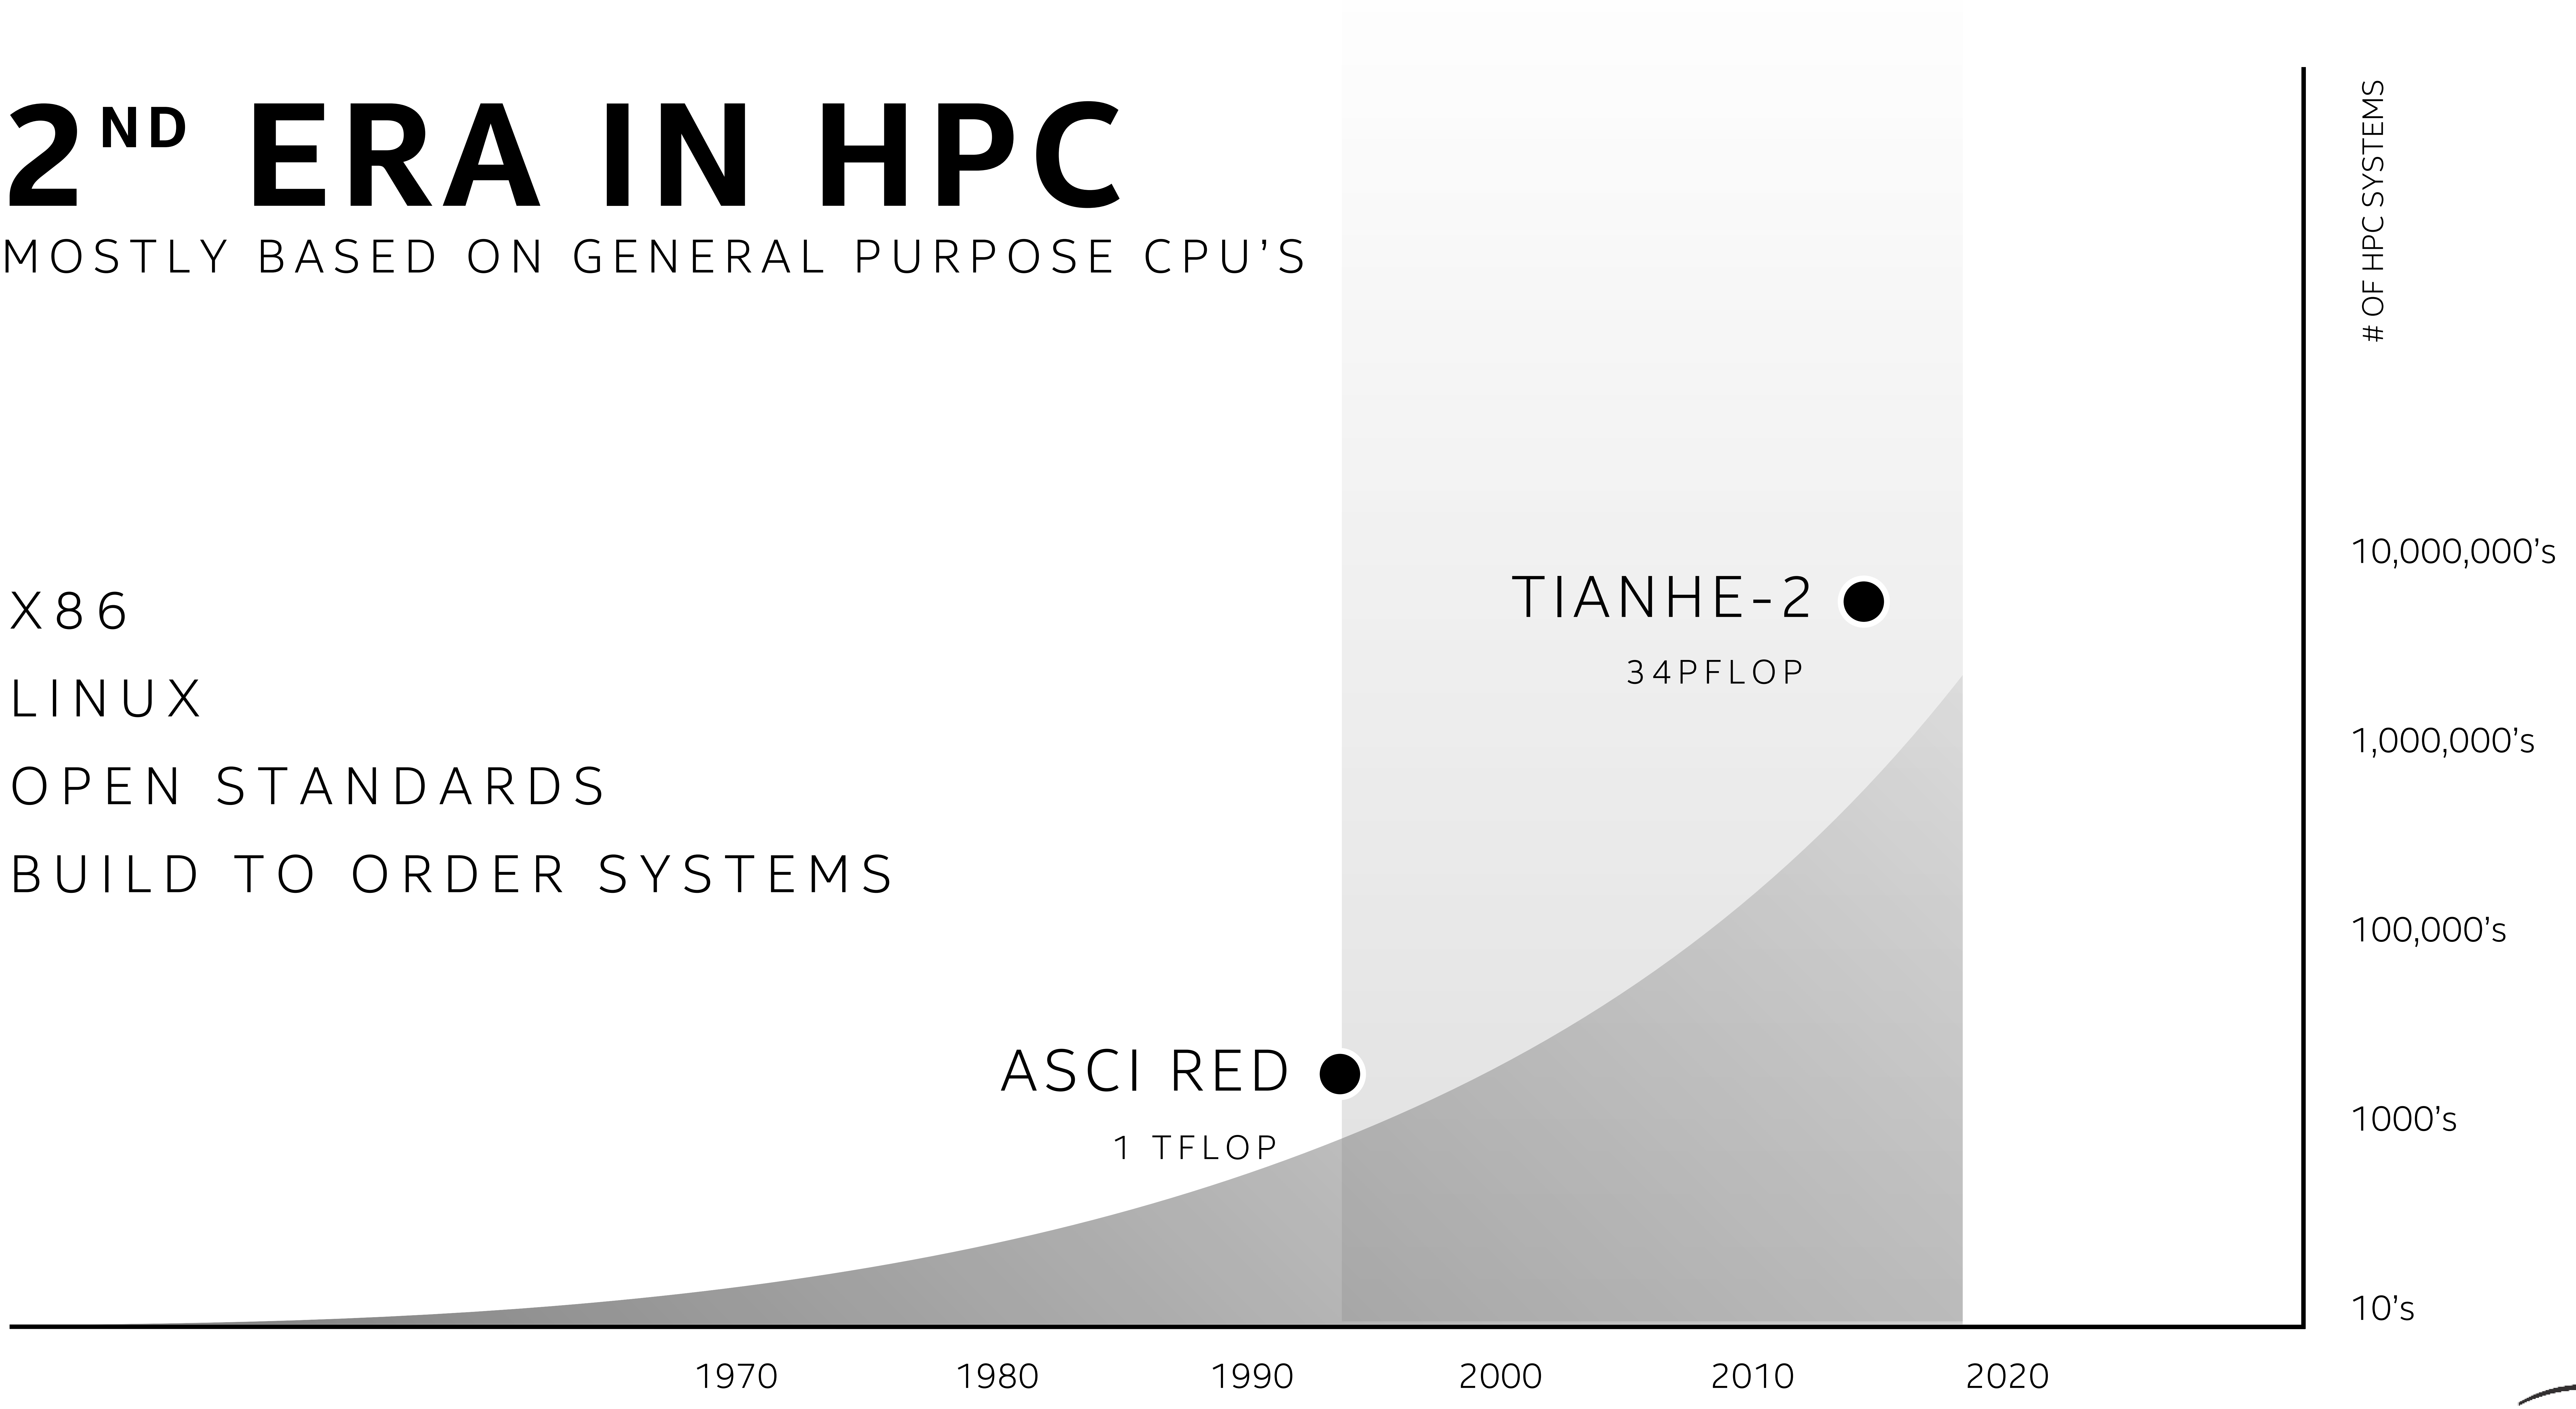

NEXT ERA IN HPC

DRIVEN BY INSATIABLE AI COMPUTE

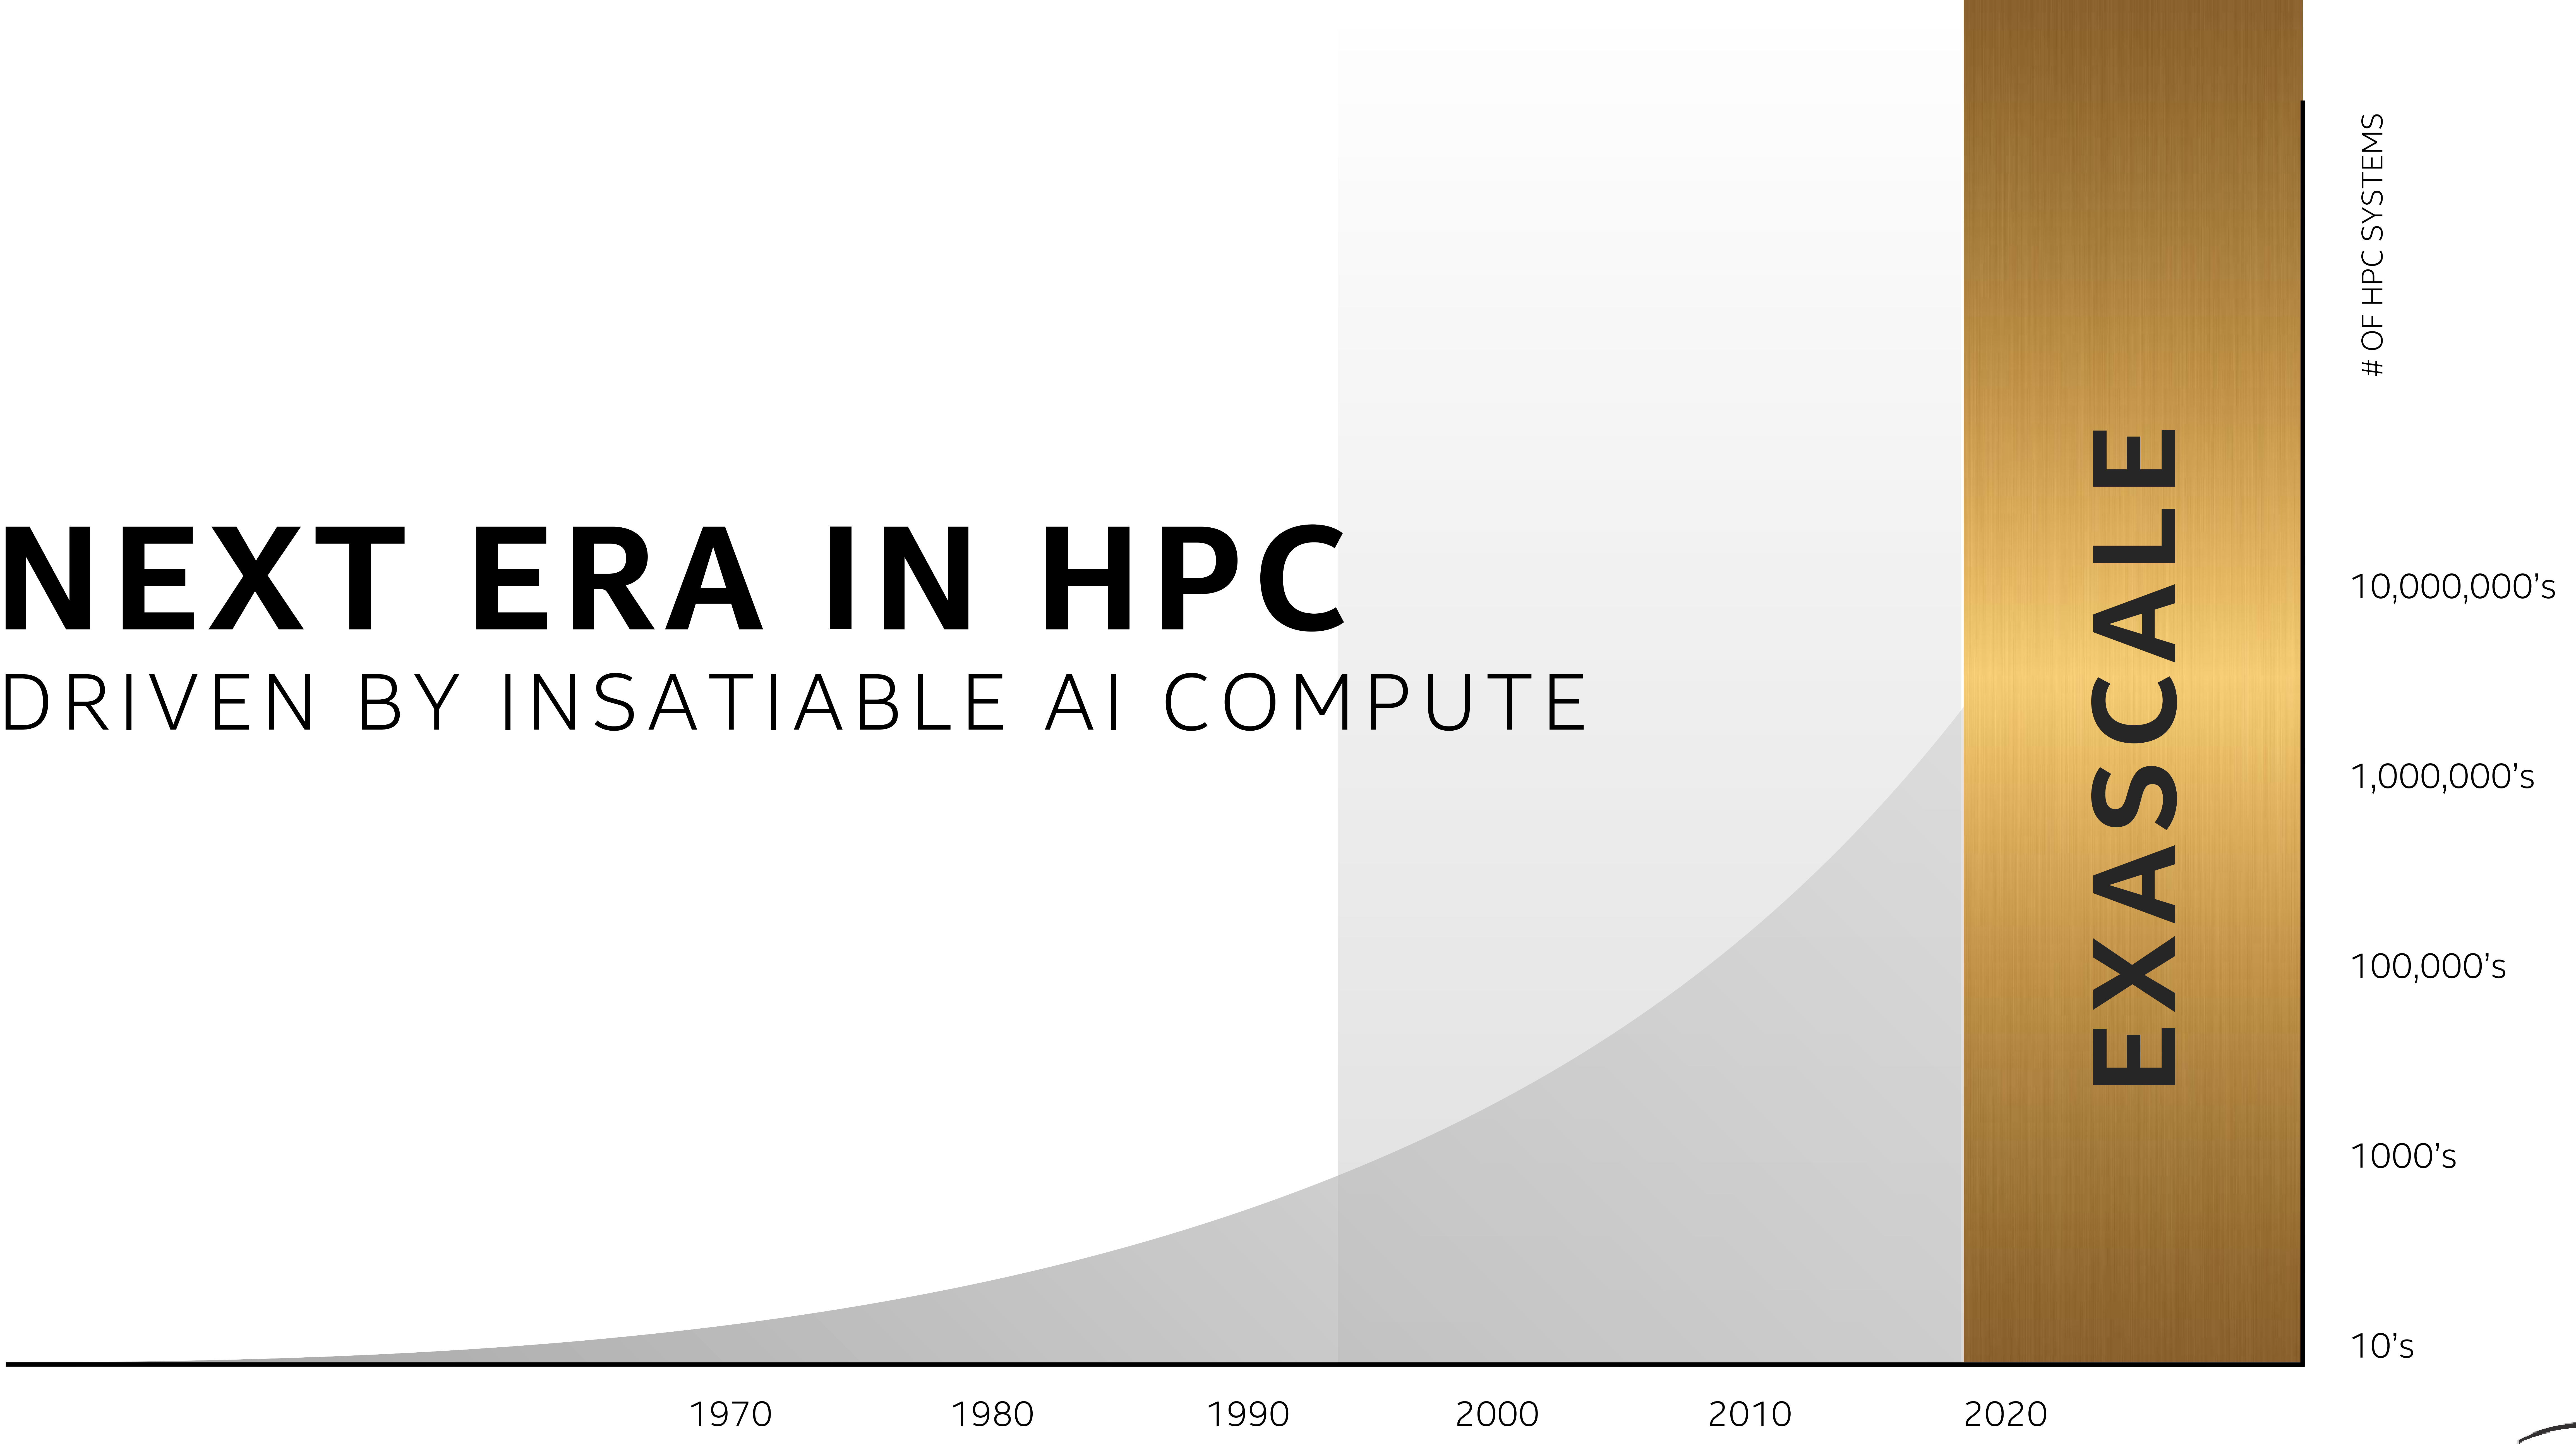

COMPUTE DEMOCRATIZATION

TECHNOLOGY LED DISRUPTIONS

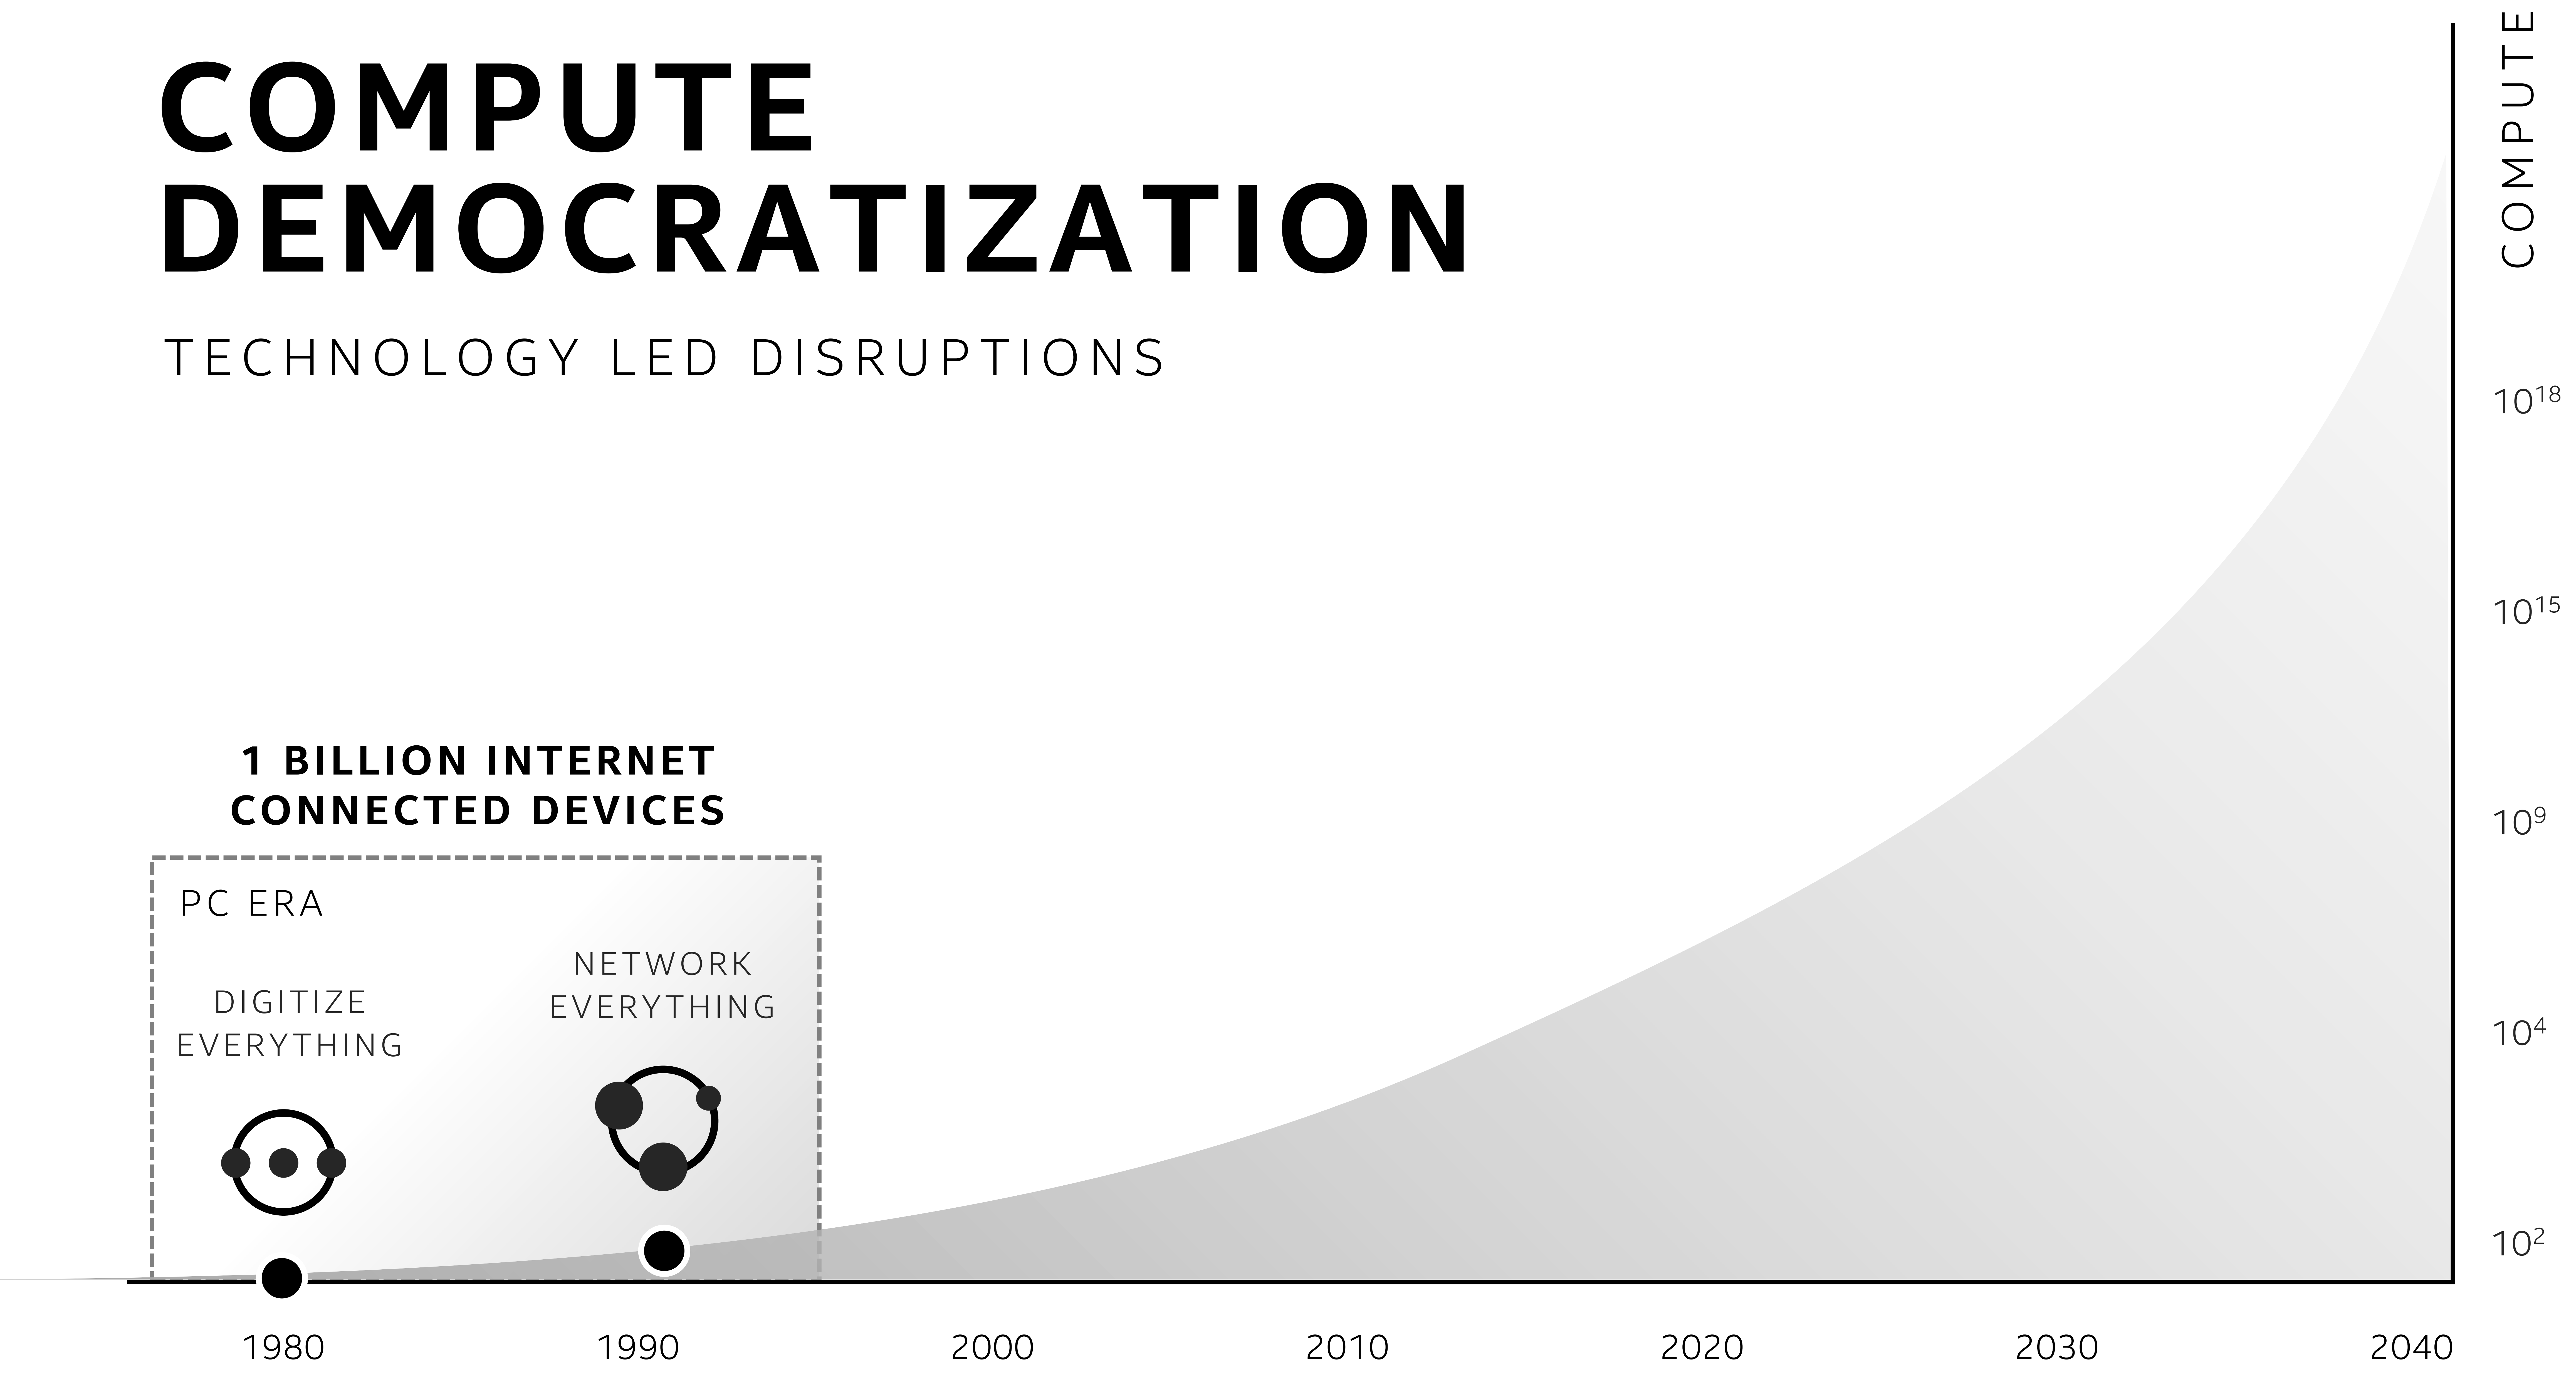

COMPUTE DEMOCRATIZATION

TECHNOLOGY LED DISRUPTIONS

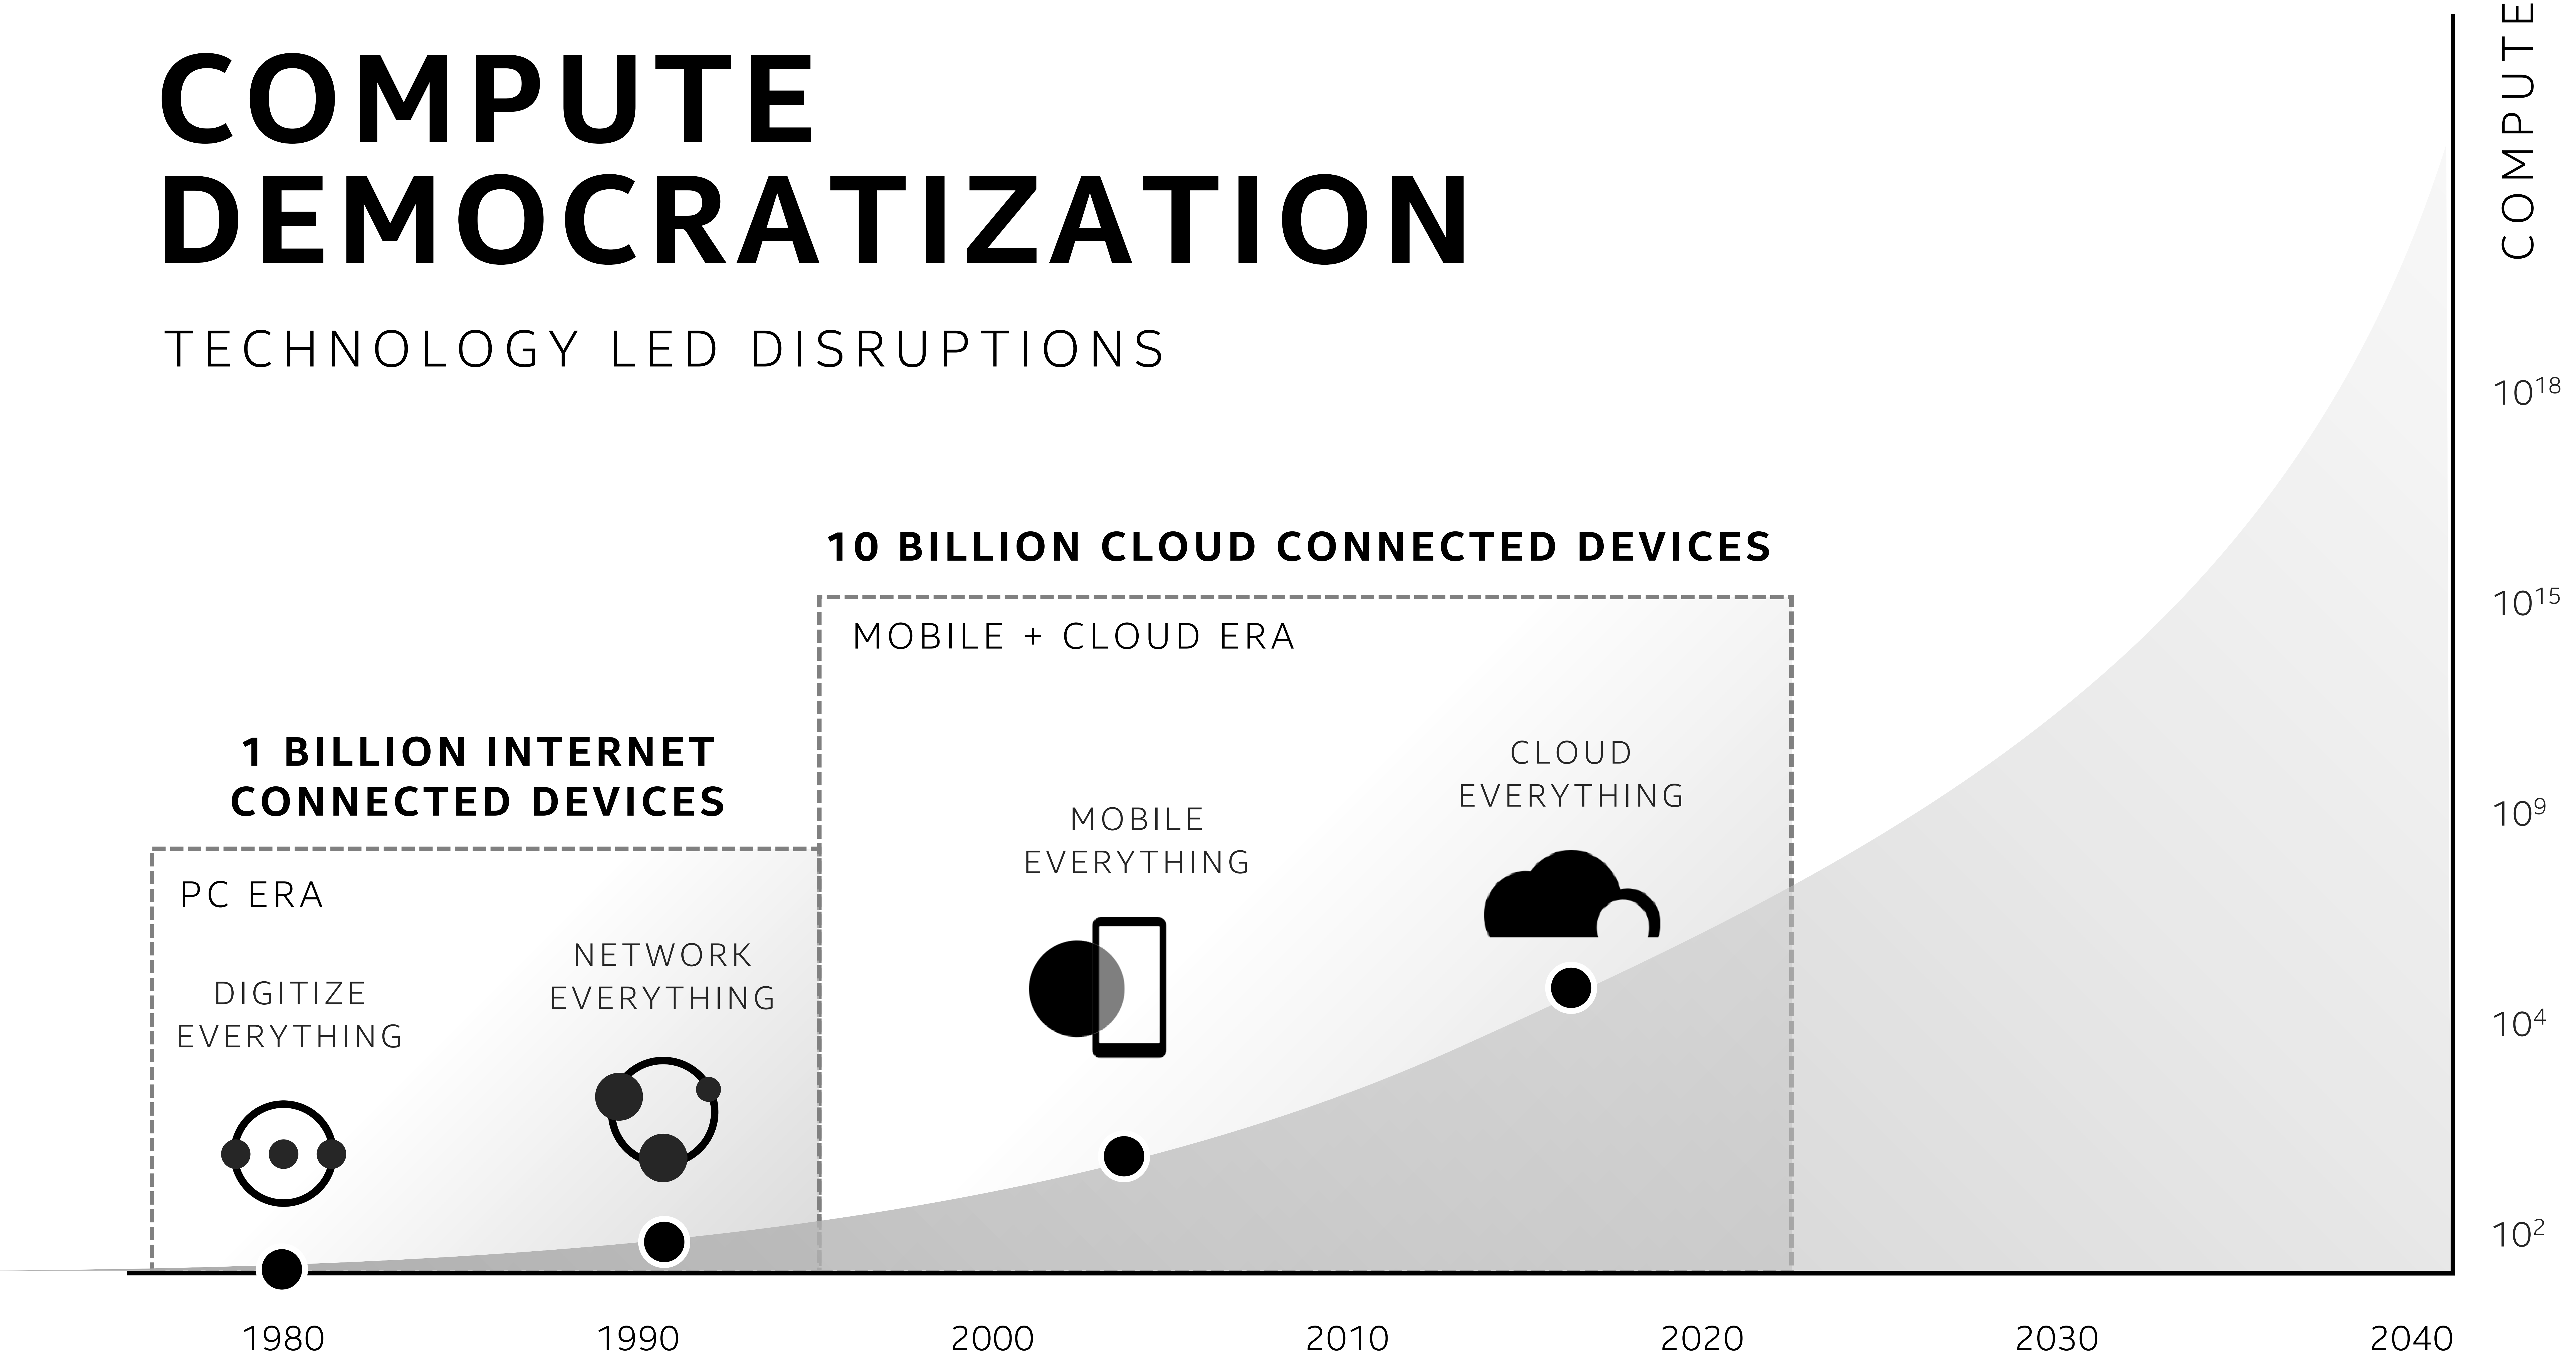

COMPUTE DEMOCRATIZATION

TECHNOLOGY LED DISRUPTIONS

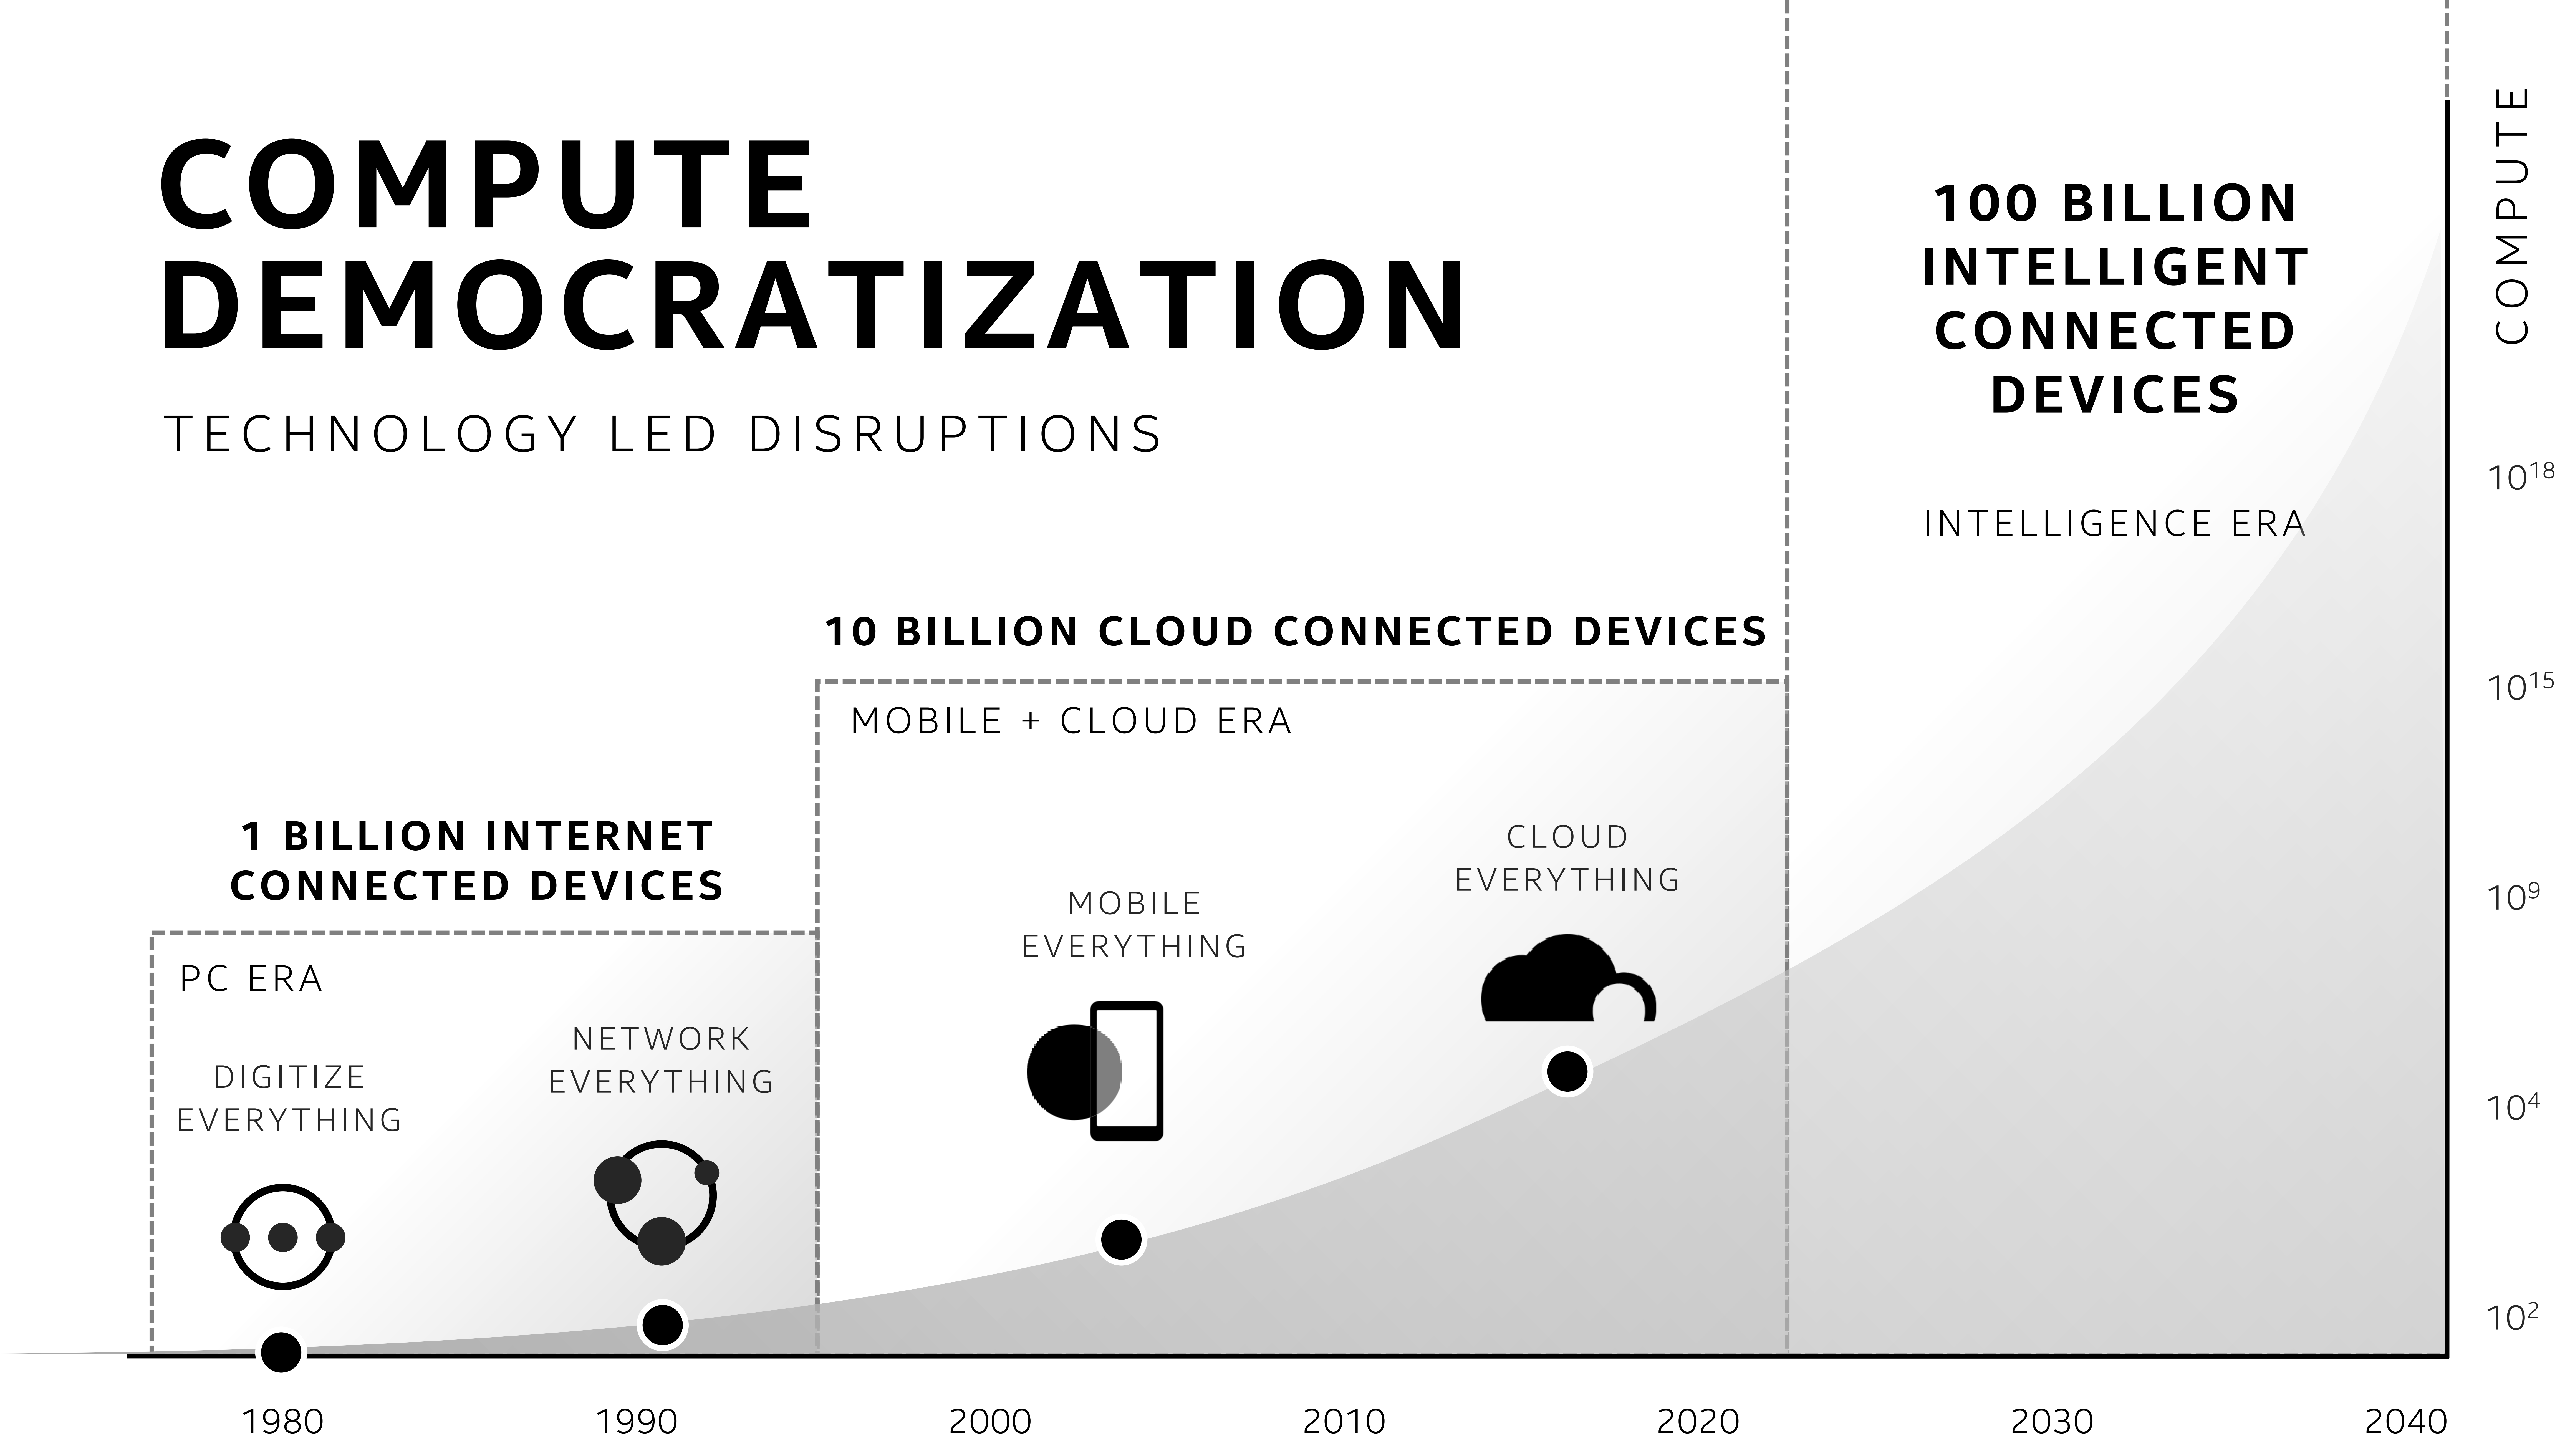


COMPUTE
DEMOCRATIZATION
TECHNOLOGY LED DISRUPTIONS
EXASCALE
FOR EVERYONE
COMPUTE
10^{18}
10^{15}
10^{9}
10^{4}
10^{2}
10 BILLION CLOUD CONNECTED DEVICES
MOBILE + CLOUD ERA
MOBILE
EVERYTHING
CLOUD
EVERYTHING
1 BILLION INTERNET
CONNECTED DEVICES
PC ERA
DIGITIZE
EVERYTHING
NETWORK
EVERYTHING
1980
1990
2000
2010
2020
2030
2040
intel

EXASCALE FOR EVERYONE

HETEROGENEOUS TAXONOMY

SCALAR | VECTOR | MATRIX | SPATIAL

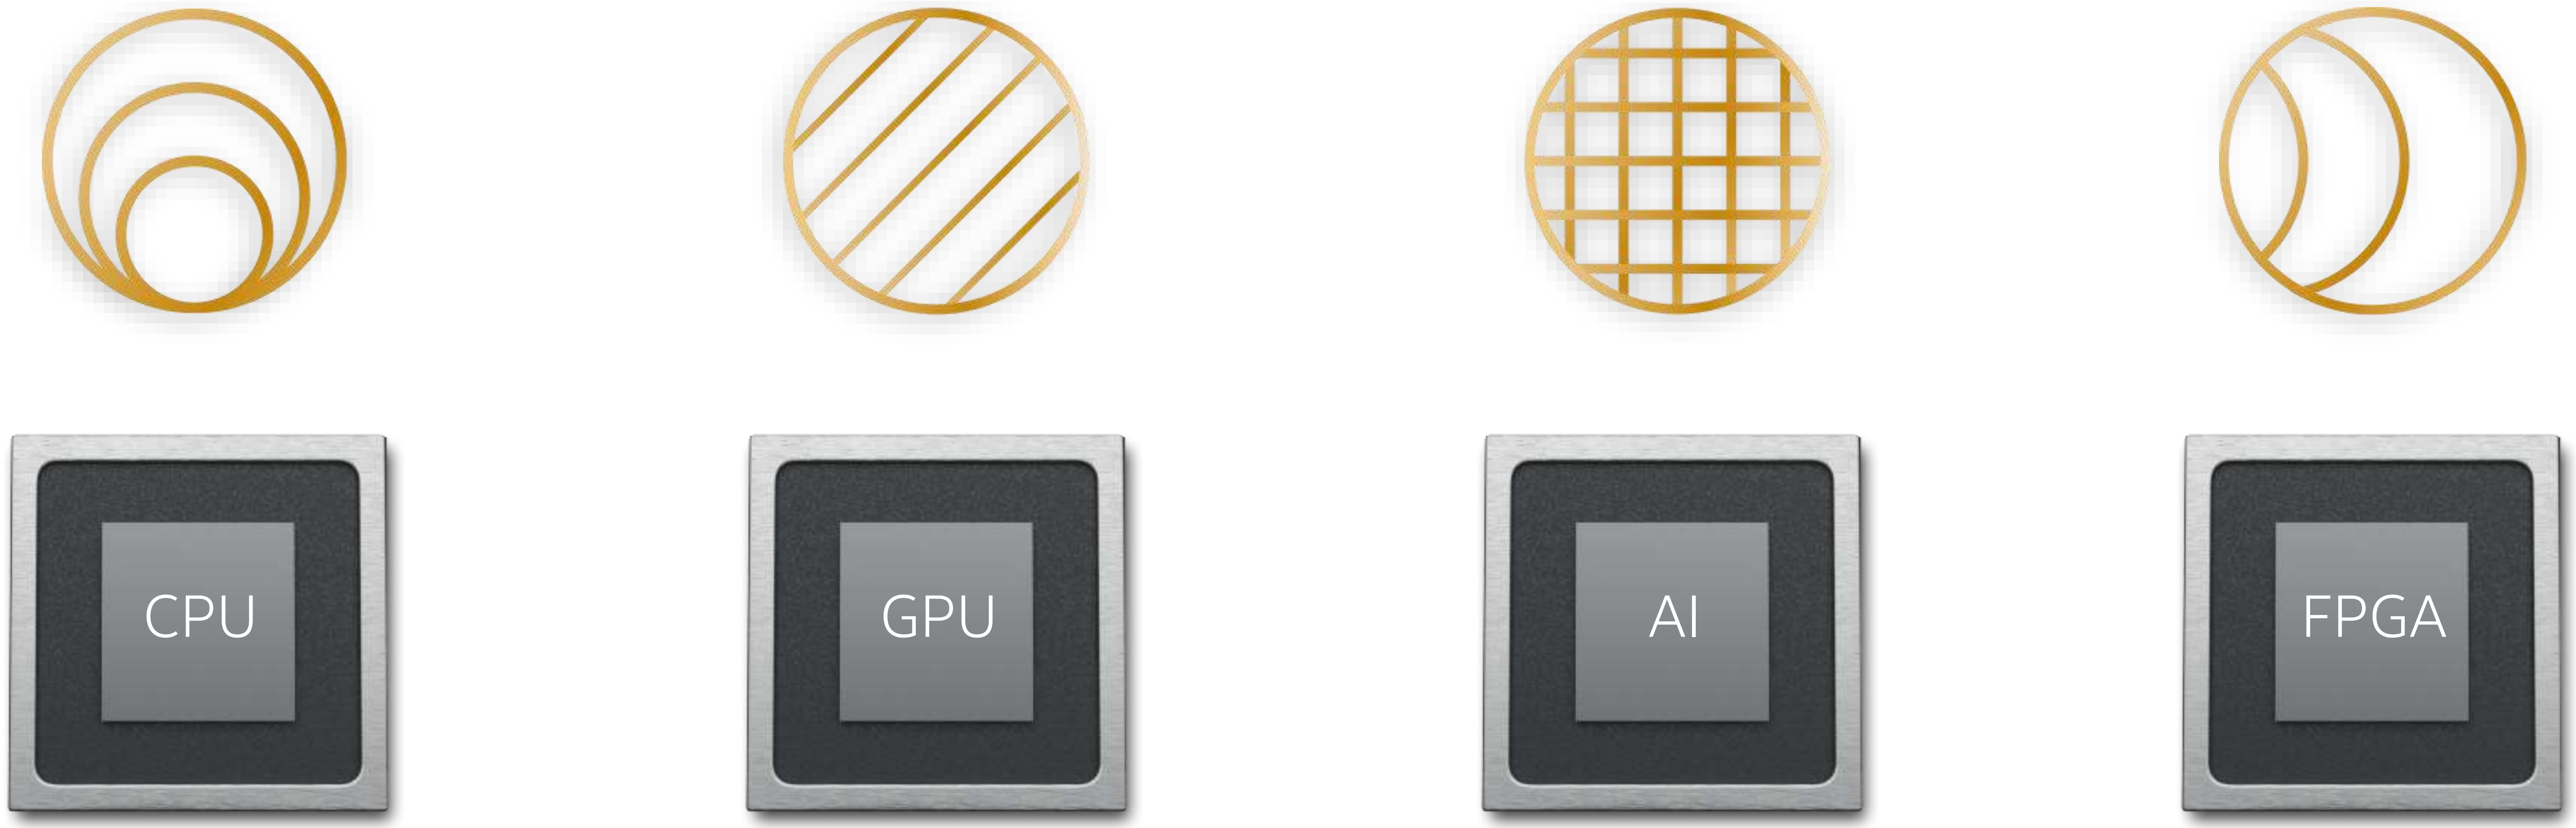

Xe HPC FEATURES

- COMPUTE
  - SCALABILITY
  - AI PERFORMANCE
  - HPC PERFORMANCE
- MEMORY
  - SCALABILITY
  - BANDWIDTH
  - UNIFIED MEMORY

TECHNOLOGY FOR EXASCALE

NEED HUGE LEAP IN PERF/WATT AND PERF/MM^2

INTEL NEXT GEN 7nm PROCESS & FOVEROS™ PACKAGING

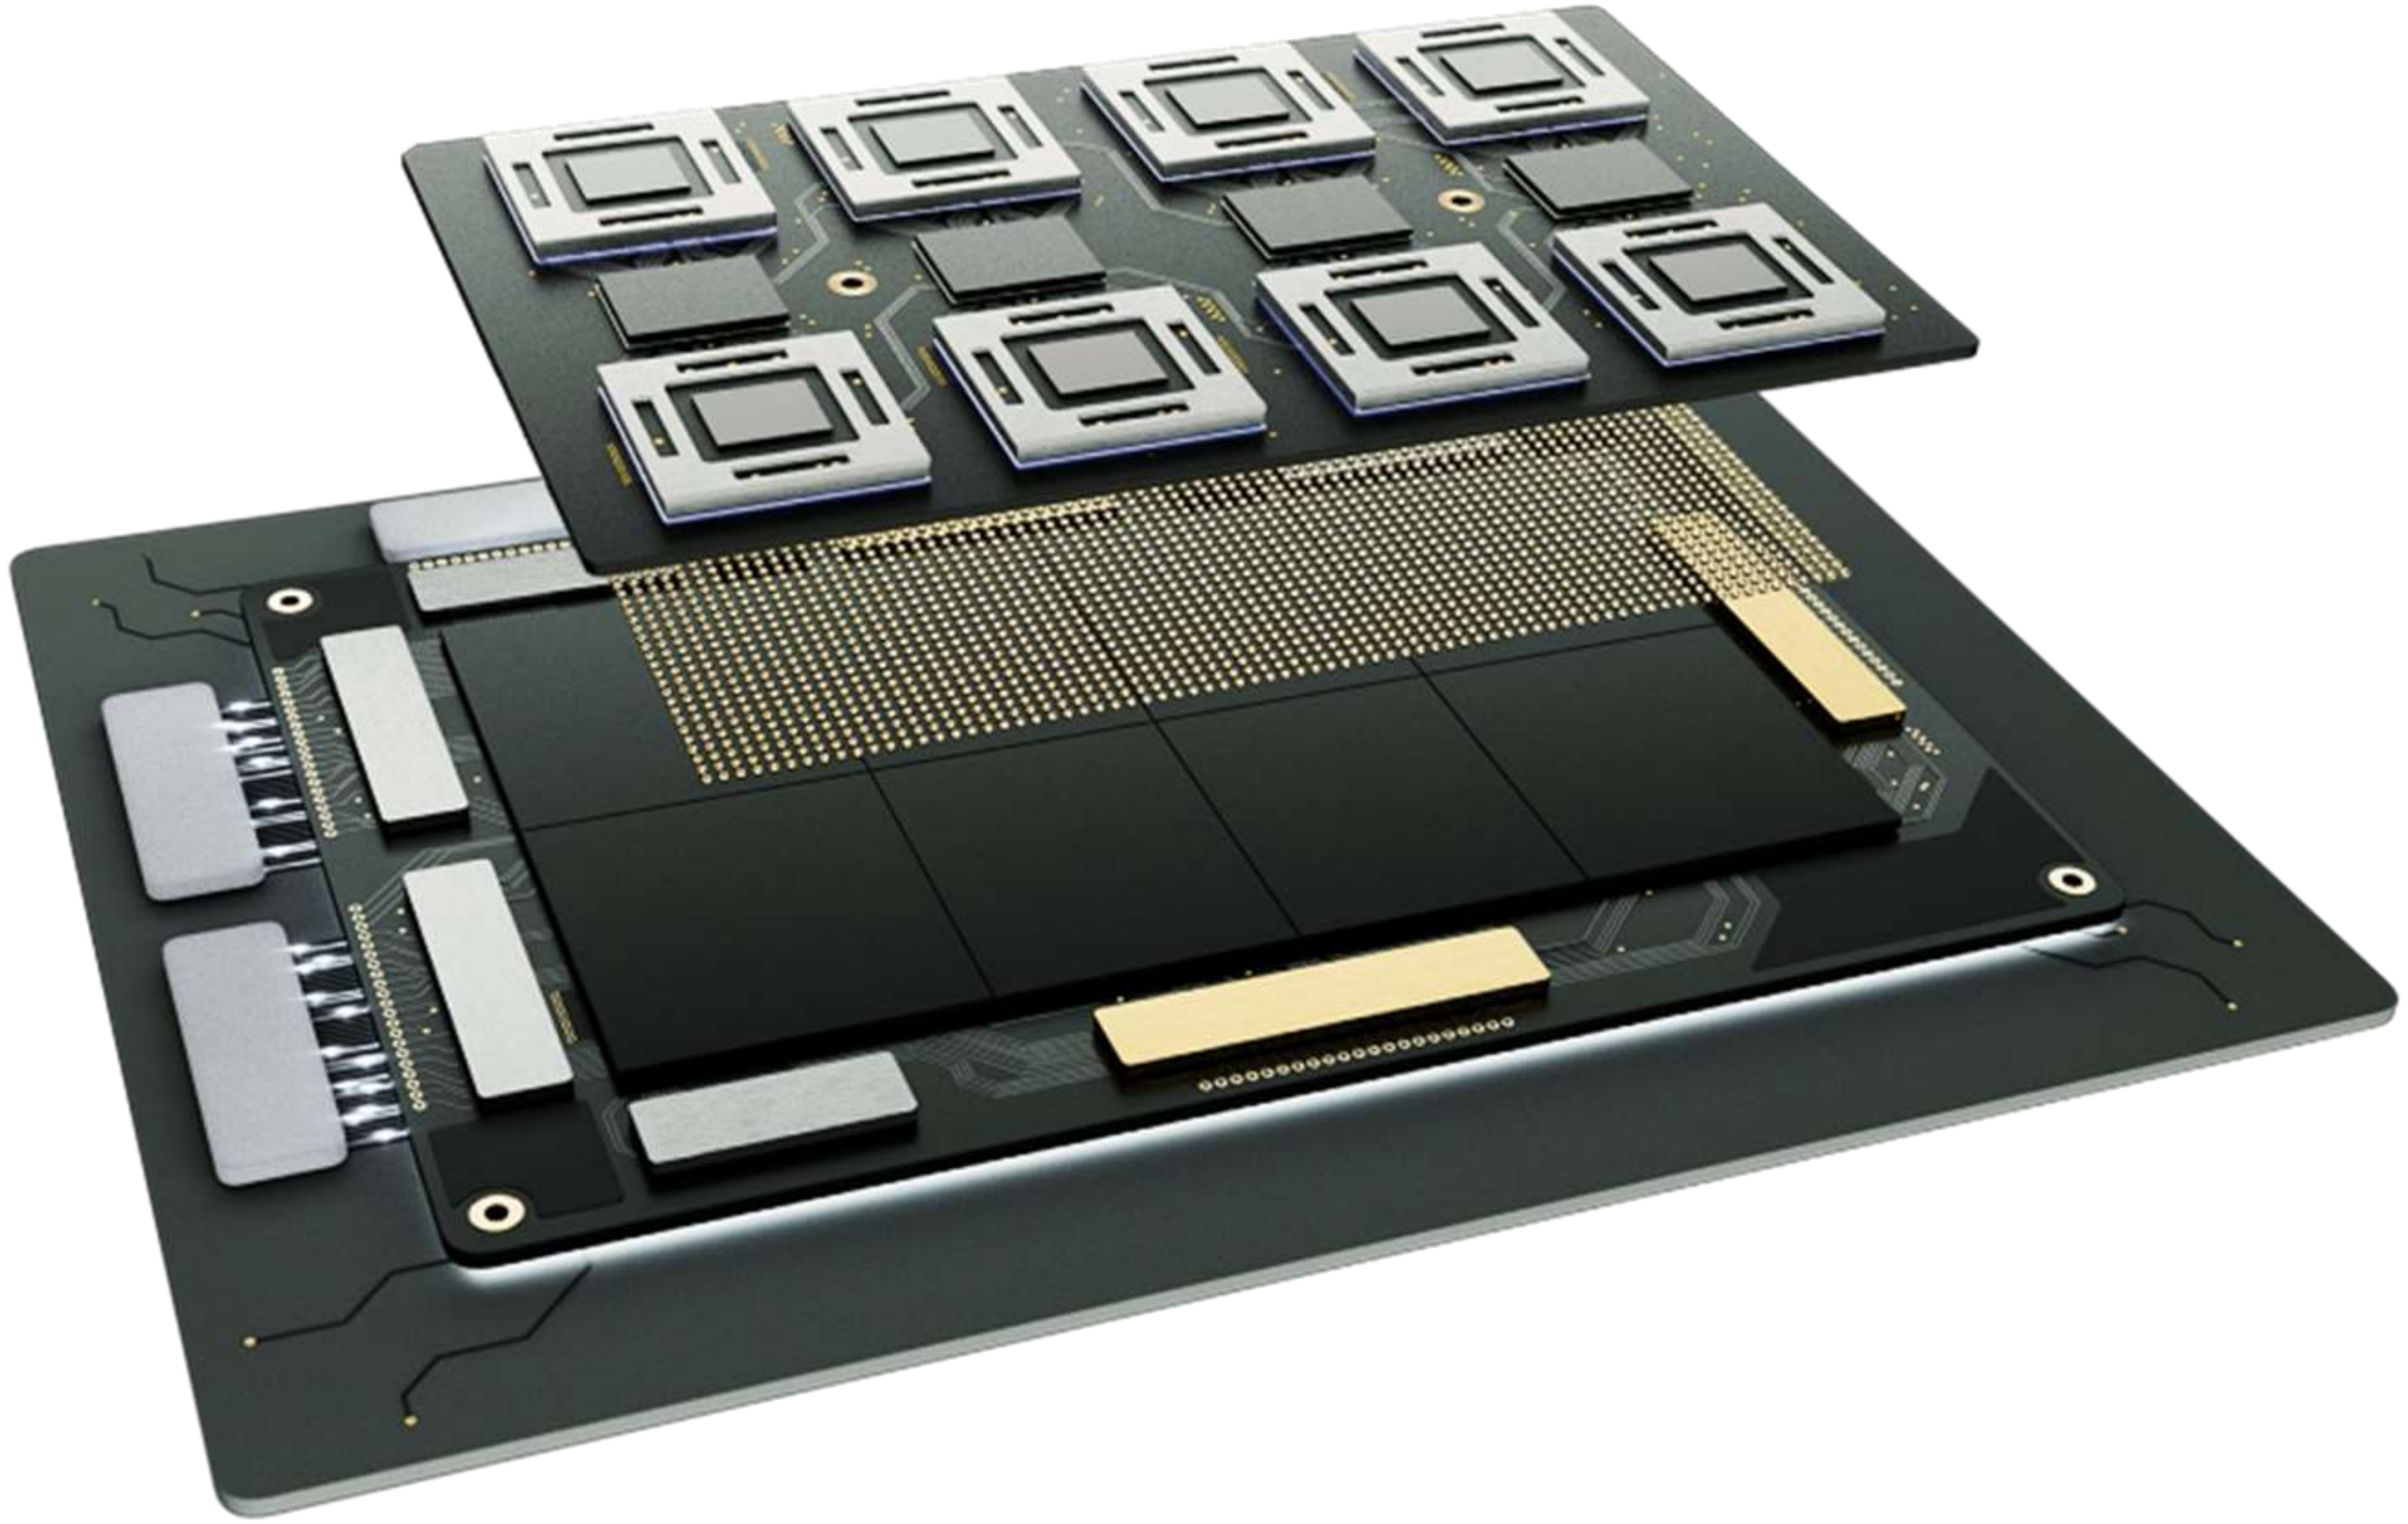

COMPUTE DENSITY

MEMORY
FOR EXASCALE

NEED HUGE LEAP IN
BANDWIDTH/WATT
& FOOTPRINT/MM^2

EMIB FOR HBM

&

FOVEROS™ FOR RAMBO CACHE

MEMORY

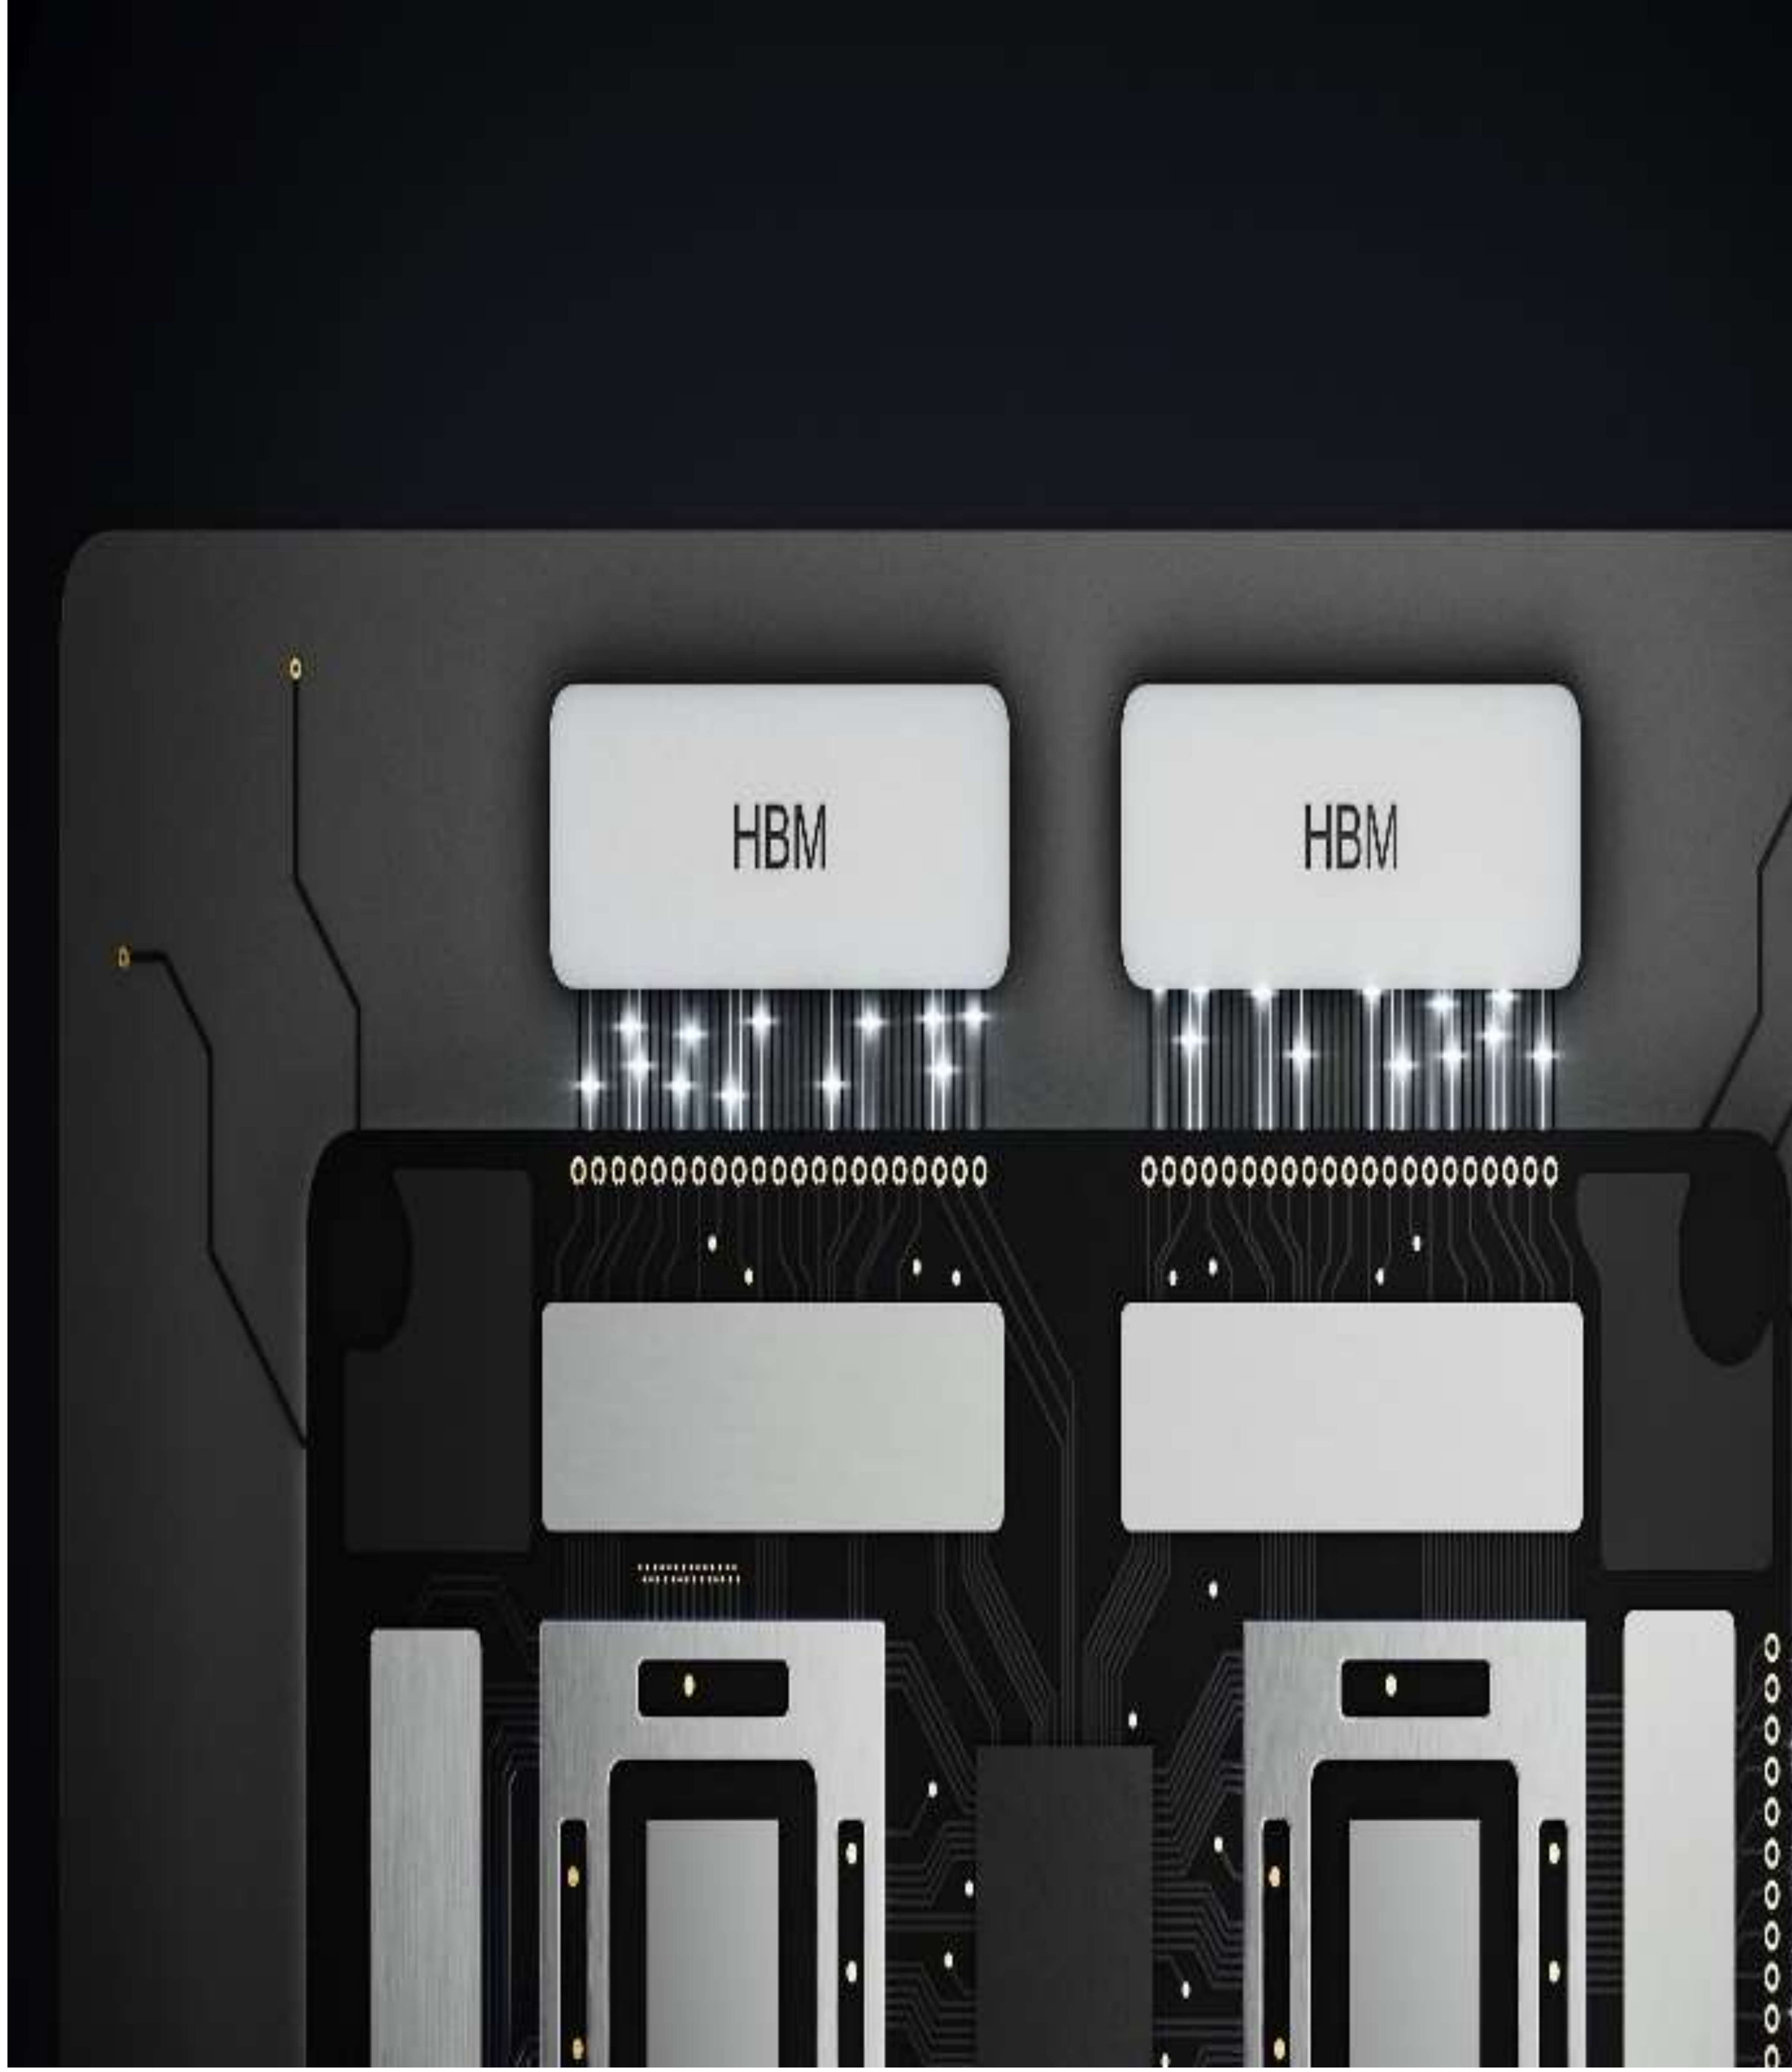

CONNECTIVITY FOR EXASCALE

Xe LINK

SCALE OUT TO MANY GPUS/NODE
UNIFIED MEMORY
CXL BASED

CONNECTIVITY

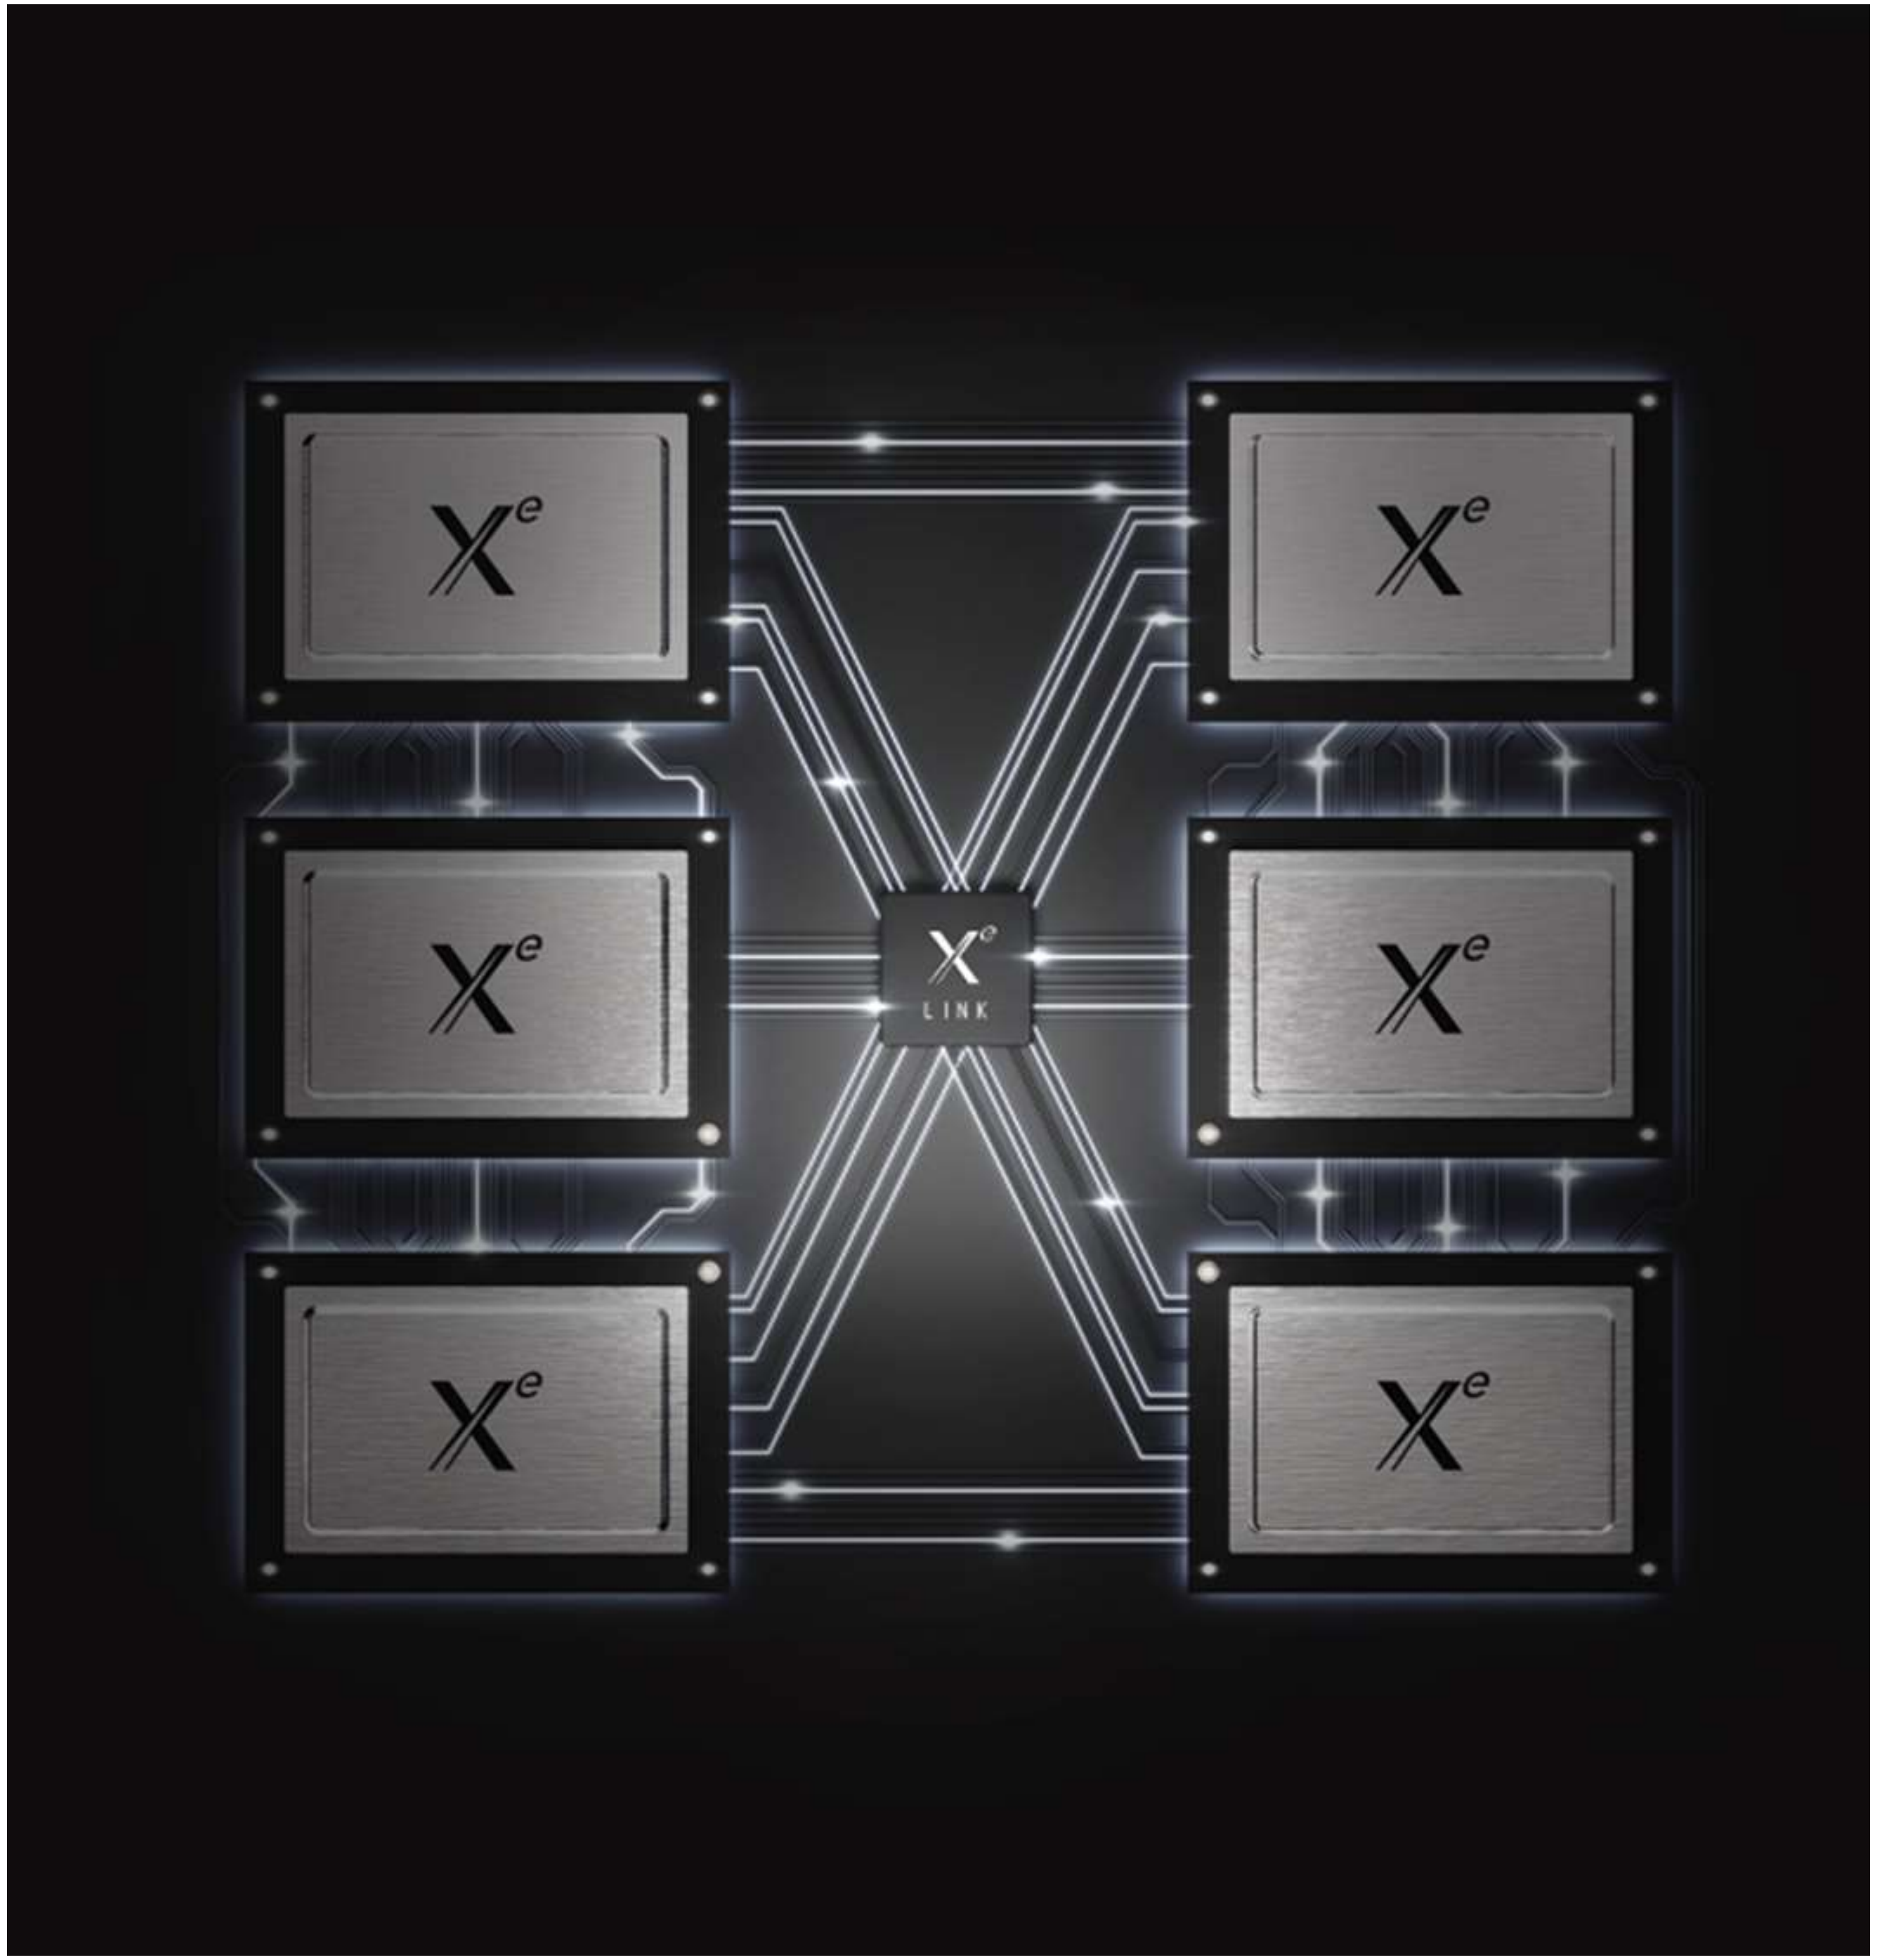

RELIABILITY FOR EXASCALE

XEON™ CLASS RAS

IN-FIELD REPAIR

ECC, PARITY ACROSS ALL MEMORY AND CACHES

CONNECTIVITY

EXASCALE GPU

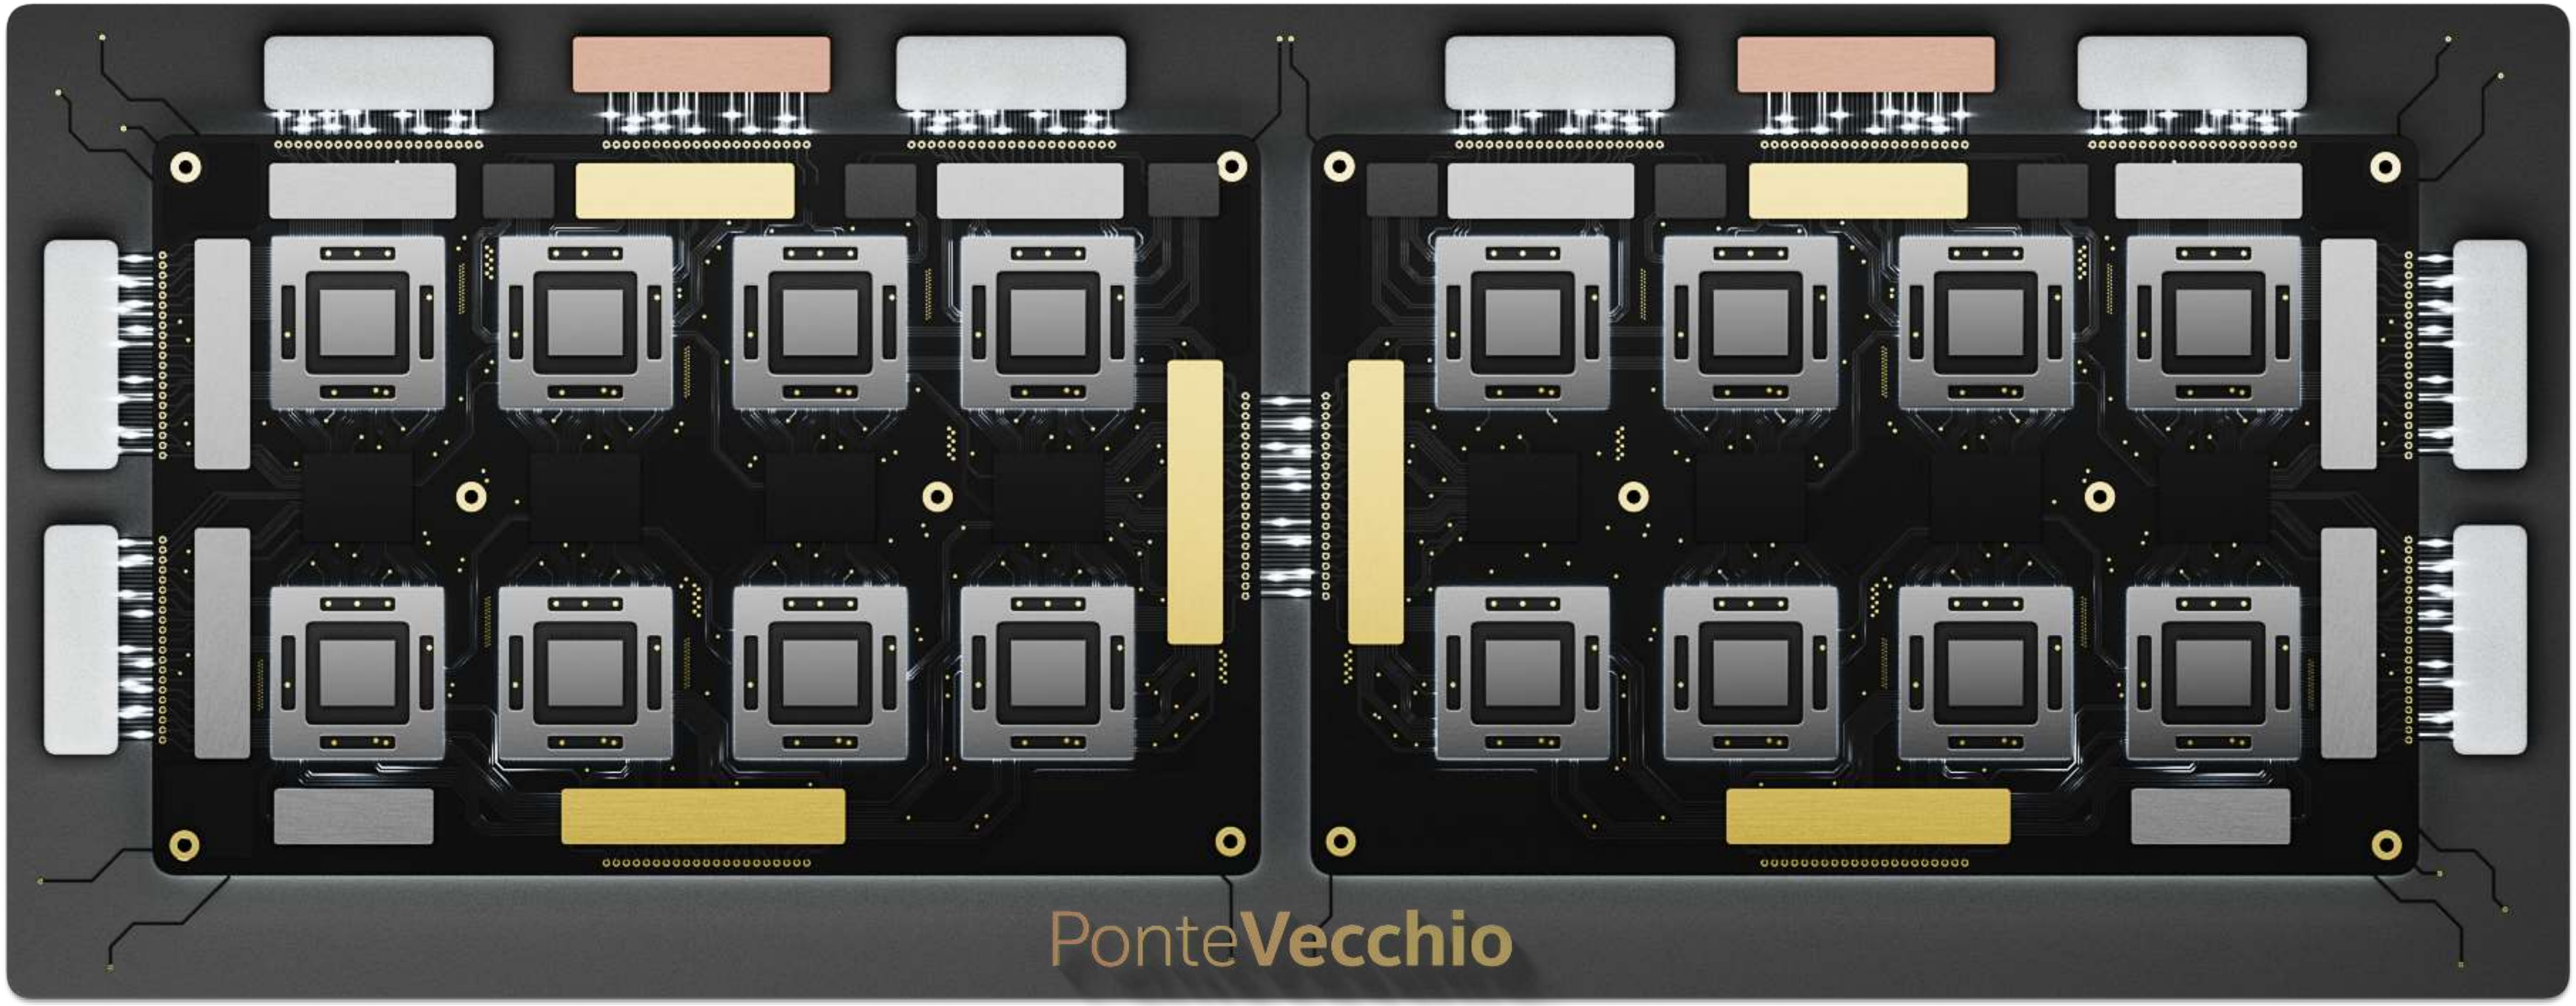

Xe HPC

intel

インテル® Xeon® プロセッサーさらなる発展

唯一、コンバージェンス向けに最適化されたデータセンター CPU

インテル® アドバンスト・ベクトル・エクステンション 512

インテル® ディープラーニング・ブースト (インテル® DL ブースト)

インテル® Optane™ DC パーシステント・メモリー

2019年	2020年	2021年
CASCADE LAKE 14NM AI の新たな加速化 (VNNI) メモリーストレージの新しい階層	**COOPER LAKE** 14NM 次世代インテル® DL ブースト (BFLOAT16) **ICE LAKE** 10NM 現在サンプルを出荷開始	**SAPPHIRE RAPIDS** 次世代テクノロジー

業界最先端のパフォーマンス

第 2 世代
インテル® Xeon®
スケーラブル・プロセッサー

インテル® Xeon® Platinum 9200 プロセッサー

新たなクラスの高度なパフォーマンス

インテル® Xeon® Platinum 8200 プロセッサー

インテル® Xeon® Gold 6200 プロセッサー

インテル® Xeon® Gold 5200 プロセッサー

インテル® Xeon® Silver 4200 プロセッサー

インテル® Xeon® Bronze 3200 プロセッサー

価値

トップクラスのワークロード・パフォーマンス

画期的なメモリー・イノベーション

人工知能アクセラレーション

ハードウェア支援型セキュリティー

俊敏性と使用率の向上

http://www.intel.co.jp/xeonscalable/

第 2 世代インテル® Xeon® スケーラブル・プロセッサーが実現する
インフラストラクチャー全体にわたる AI 対応

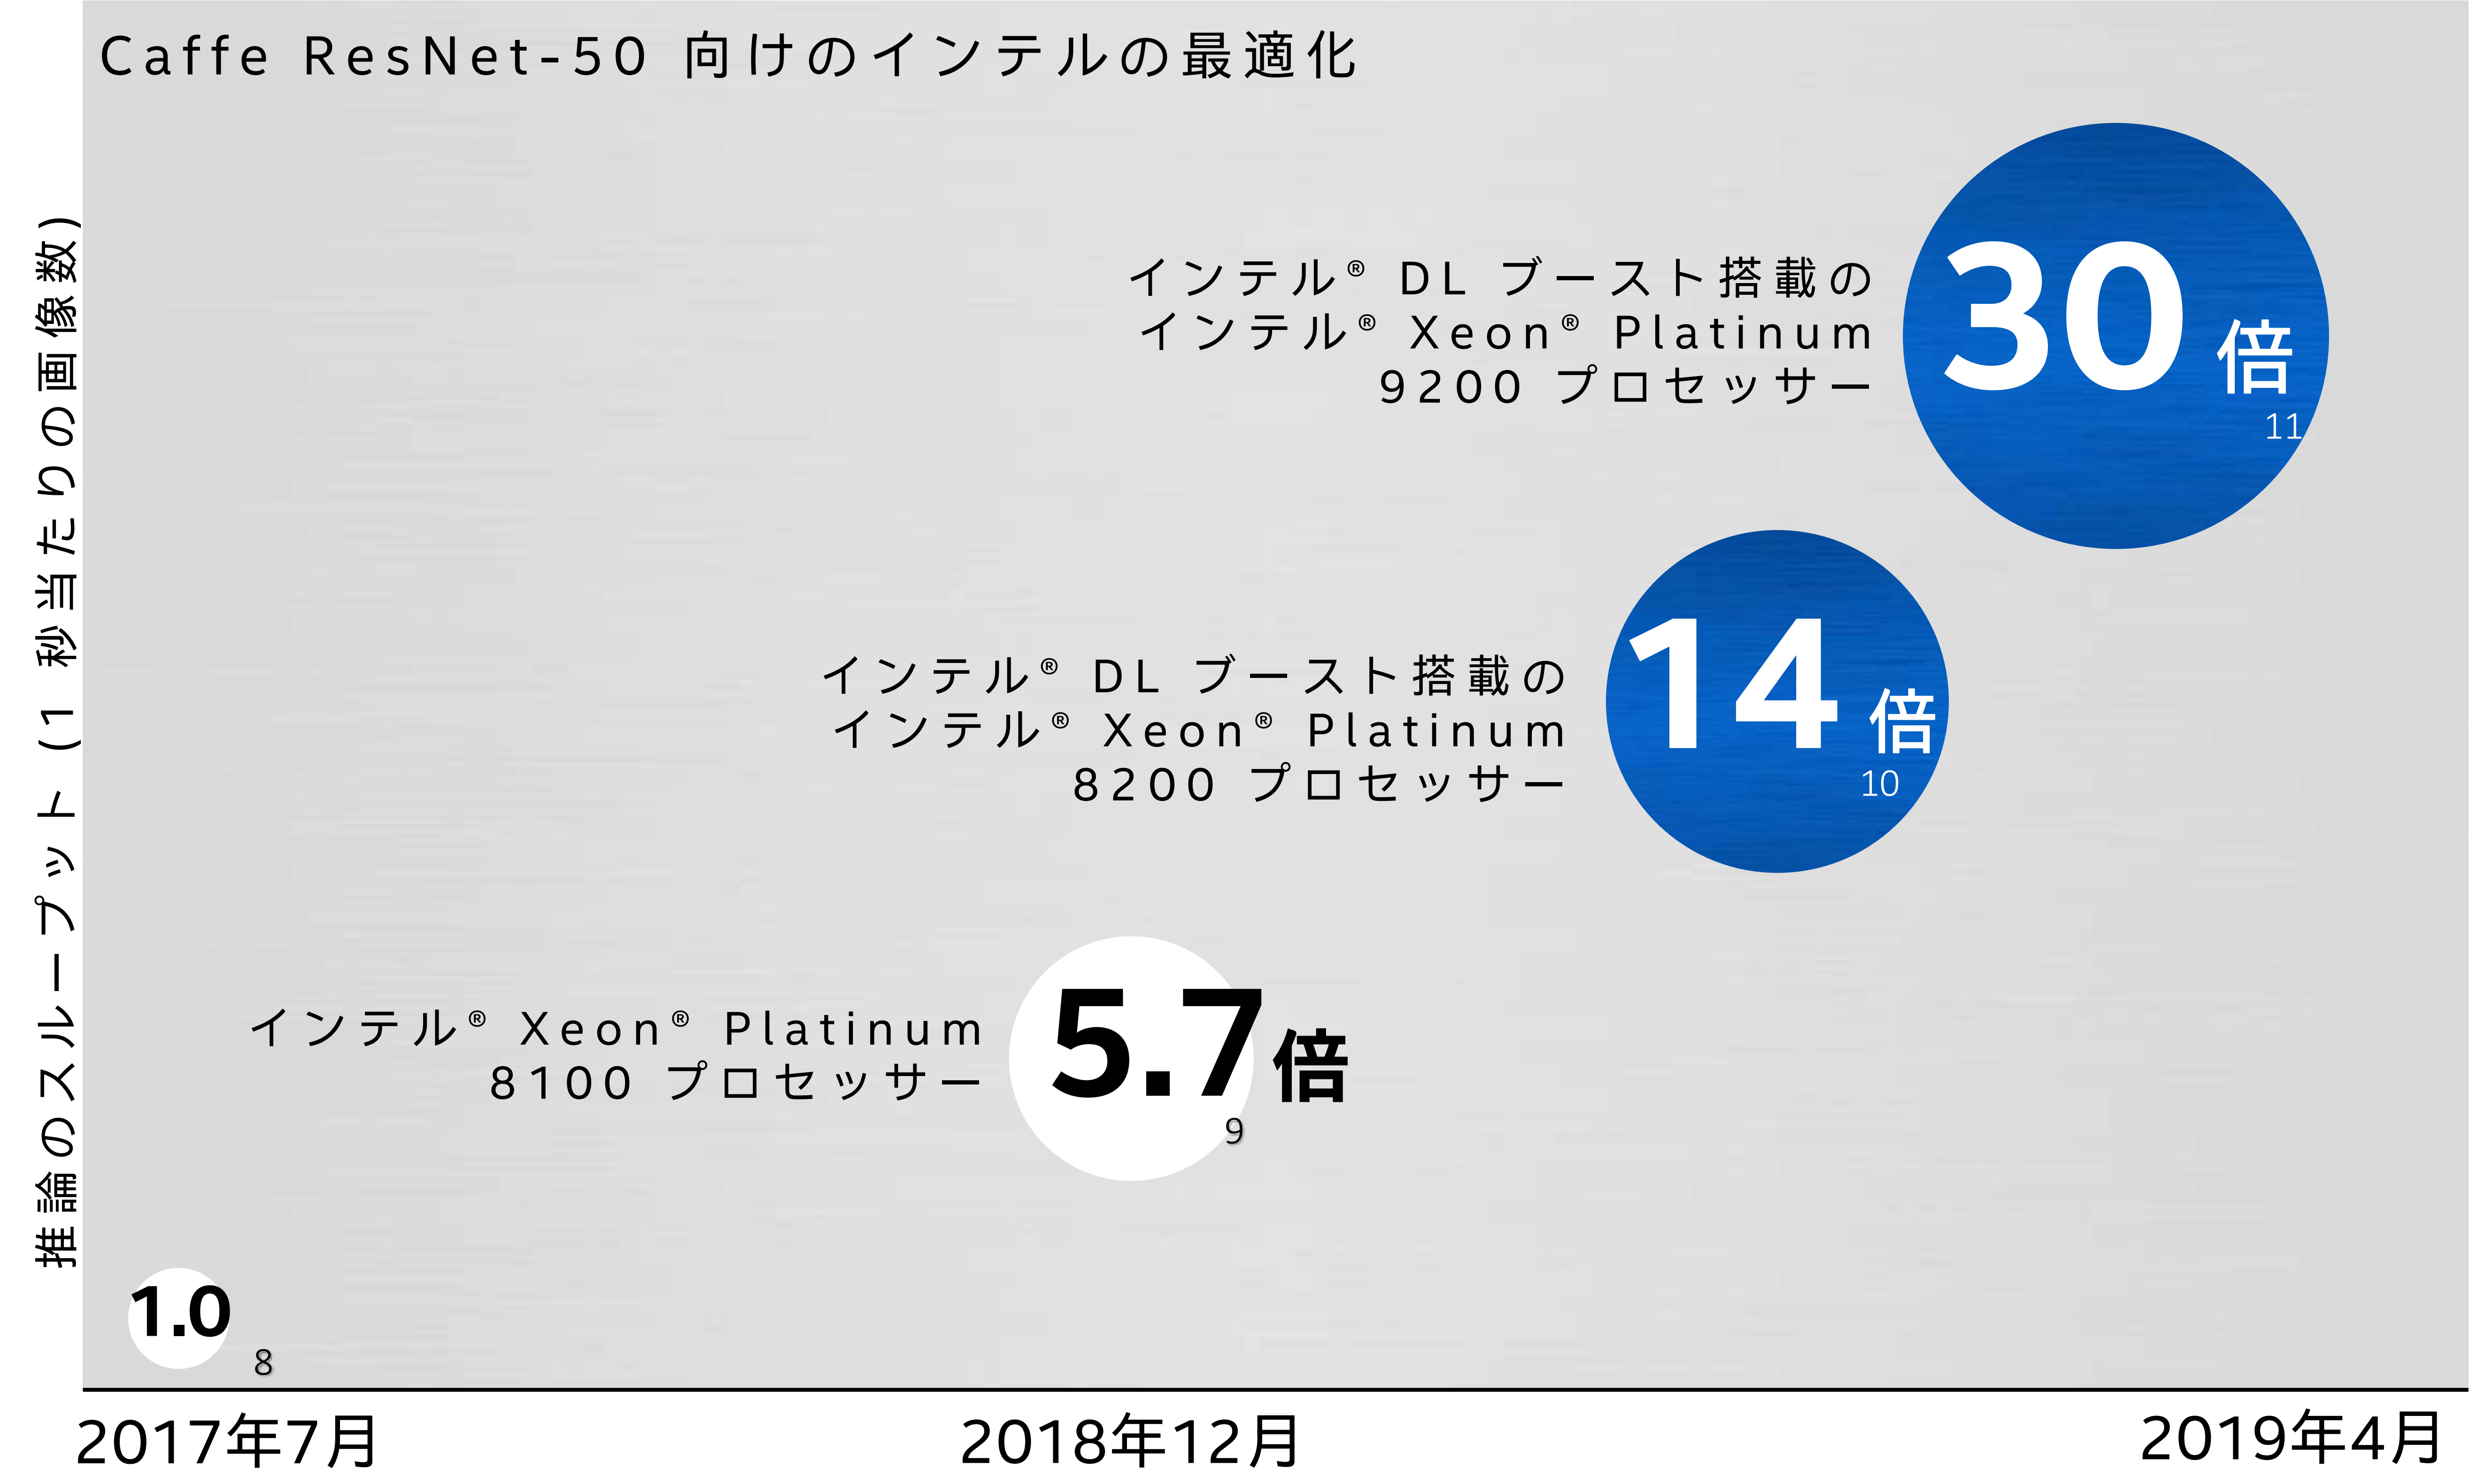

新しい AI アクセラレーション

インテル® ディープラーニング・ブースト

VNNL Vector Neural Network Instruction

内蔵の推論アクセラレーション

開発者とデータ・サイエンティスト向けの最適化

OpenVINO™ ツールキット

最適化したフレームワーク: TensorFlow, mxnet, Caffe, PYTORCH, PaddlePaddle

データタイプ: ONNX

性能の測定結果は、構成に示した日付時点のテストに基づいています。また、現在公開中のすべてのセキュリティー・アップデートが適用されているとは限りません。構成とベンチマークの詳細は、スライド 52 ページに記載しています。 絶対的なセキュリティーを提供できる製品やコンポーネントはありません。 性能に関するテストに使用されるソフトウェアとワークロードは、性能がインテル® マイクロプロセッサー用に最適化されていることがあります。SYSmark* や MobileMark* などの性能テストは、特定のコンピューター・システム、コンポーネント、ソフトウェア、操作、機能に基づいて行ったものです。結果はこれらの要因によって異なります。製品の購入を検討される場合は、他の製品と組み合わせた場合の本製品の性能など、ほかの情報や性能テストも参考にして、パフォーマンスを総合的に評価することをお勧めします。 詳細については、http://www.intel.com/benchmarks/ (英語) を参照してください。

http://ai.intel.com/ (英語)

新しいレベルのパフォーマンス

インテル® Xeon® Platinum9200 プロセッサー

最大
パフォーマンス

ソケットあたり
最大
56 コア

ラックあたりのパフォーマンス
リーダーシップ

高帯域幅

最大
12 チャンネル
ネイティブ DDR4 メモリー

データ集約度の高いワークロード用に設計

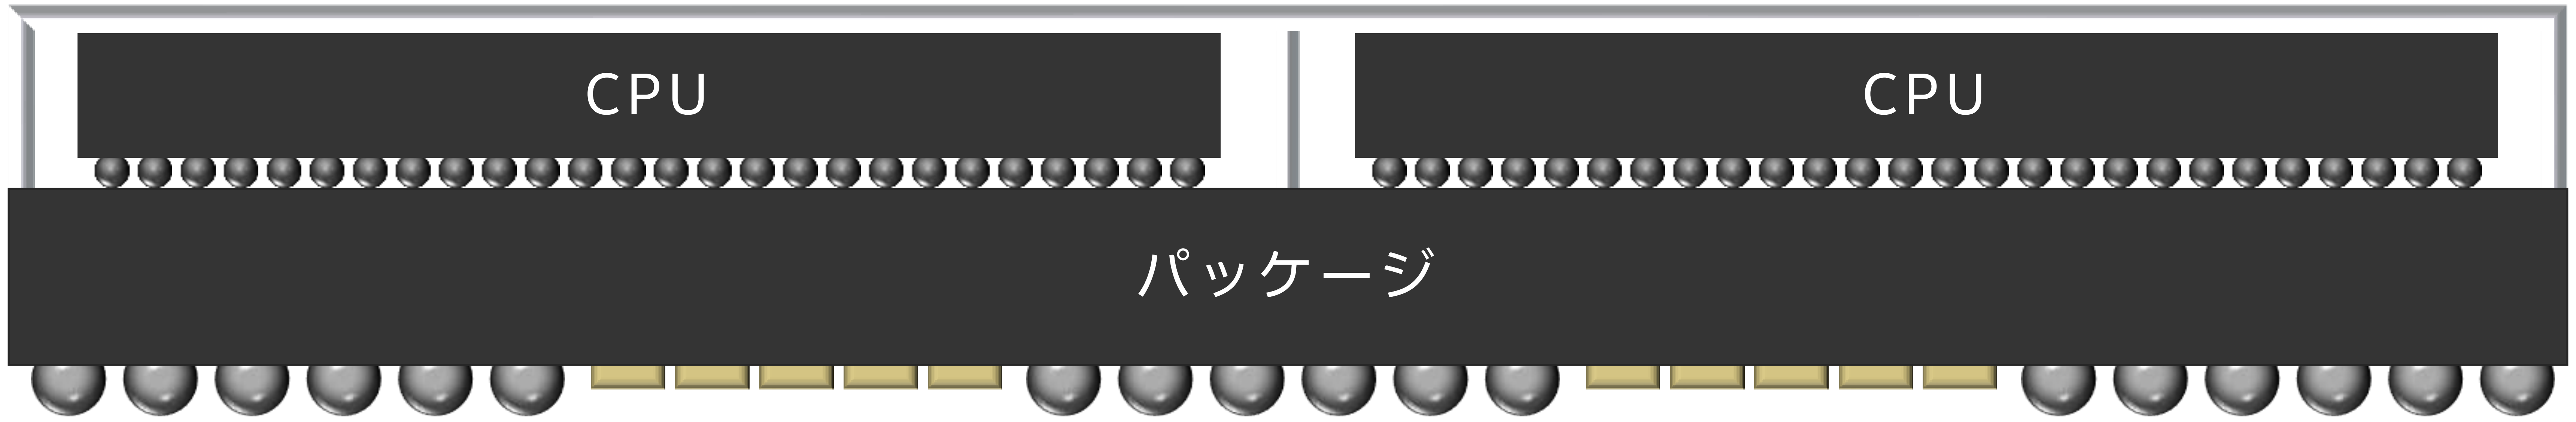

Leadership performance per rack based on 4 benchmarks (Integer Throughput, Floating Point Throughput, Memory Bandwidth and LINPACK).

intel

インテル® oneAPI

- インテル® oneAPI プロジェクトはさまざまなアーキテクチャーにわたる開発を容易にする統合プログラミング・モデルを提供
- データ並列 C++ 言語とインテル® oneAPI ライブラリー API により並列処理を表現
- 妥協のないパフォーマンス
- CPU、GPU、AI、および FPGA をサポート
- 業界標準およびオープン仕様ベース
- OpenMP*、Fortran、MPI などとの相互運用性

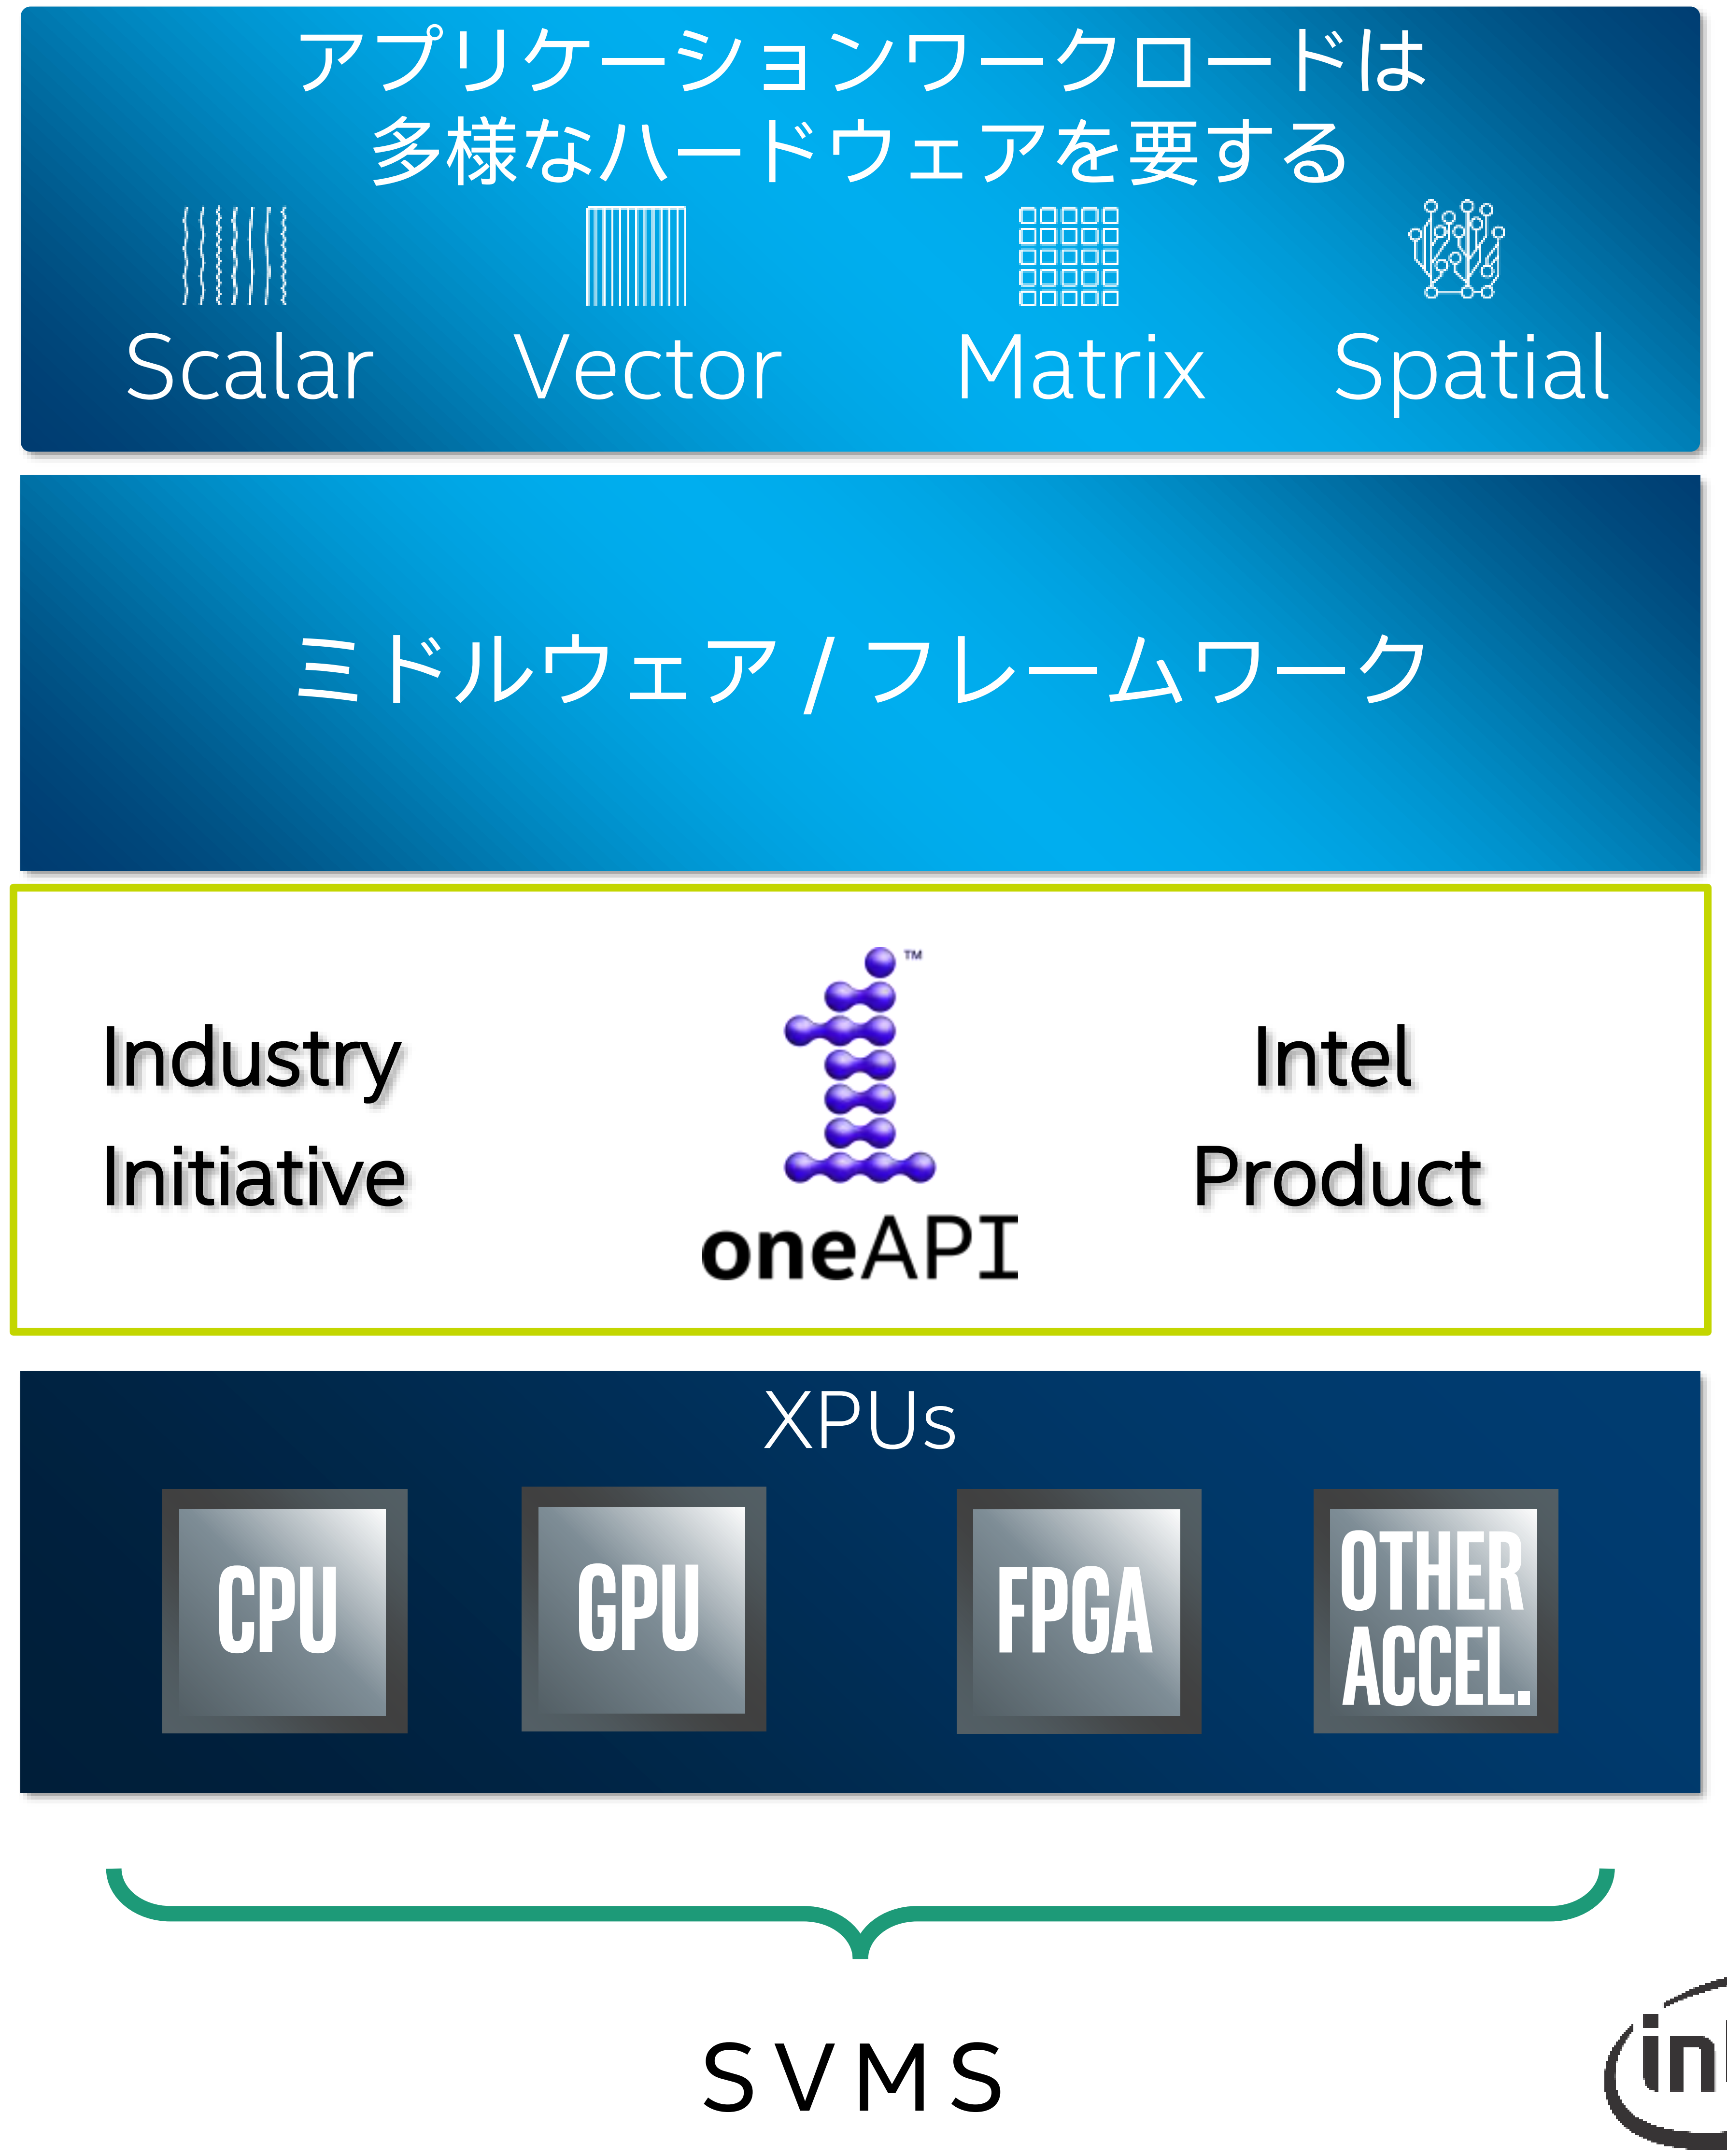

oneAPI の目指すもの：
単一ベンダーソリューションの代替

- 標準ベースのクロスアーキテクチャ言語 DPC++、C++ と SYCLがベース。
- キードメイン特化の機能を加速するために設計された強力なAPI群
- ハードウェアの抽象レイヤをベンダーへ提供する低レベルハードウェアインターフェース
- オープンスタンダード、コミュニティへの訴求、インダストリサポート
- アーキテクチャとベンダーをまたがったコードの再利用

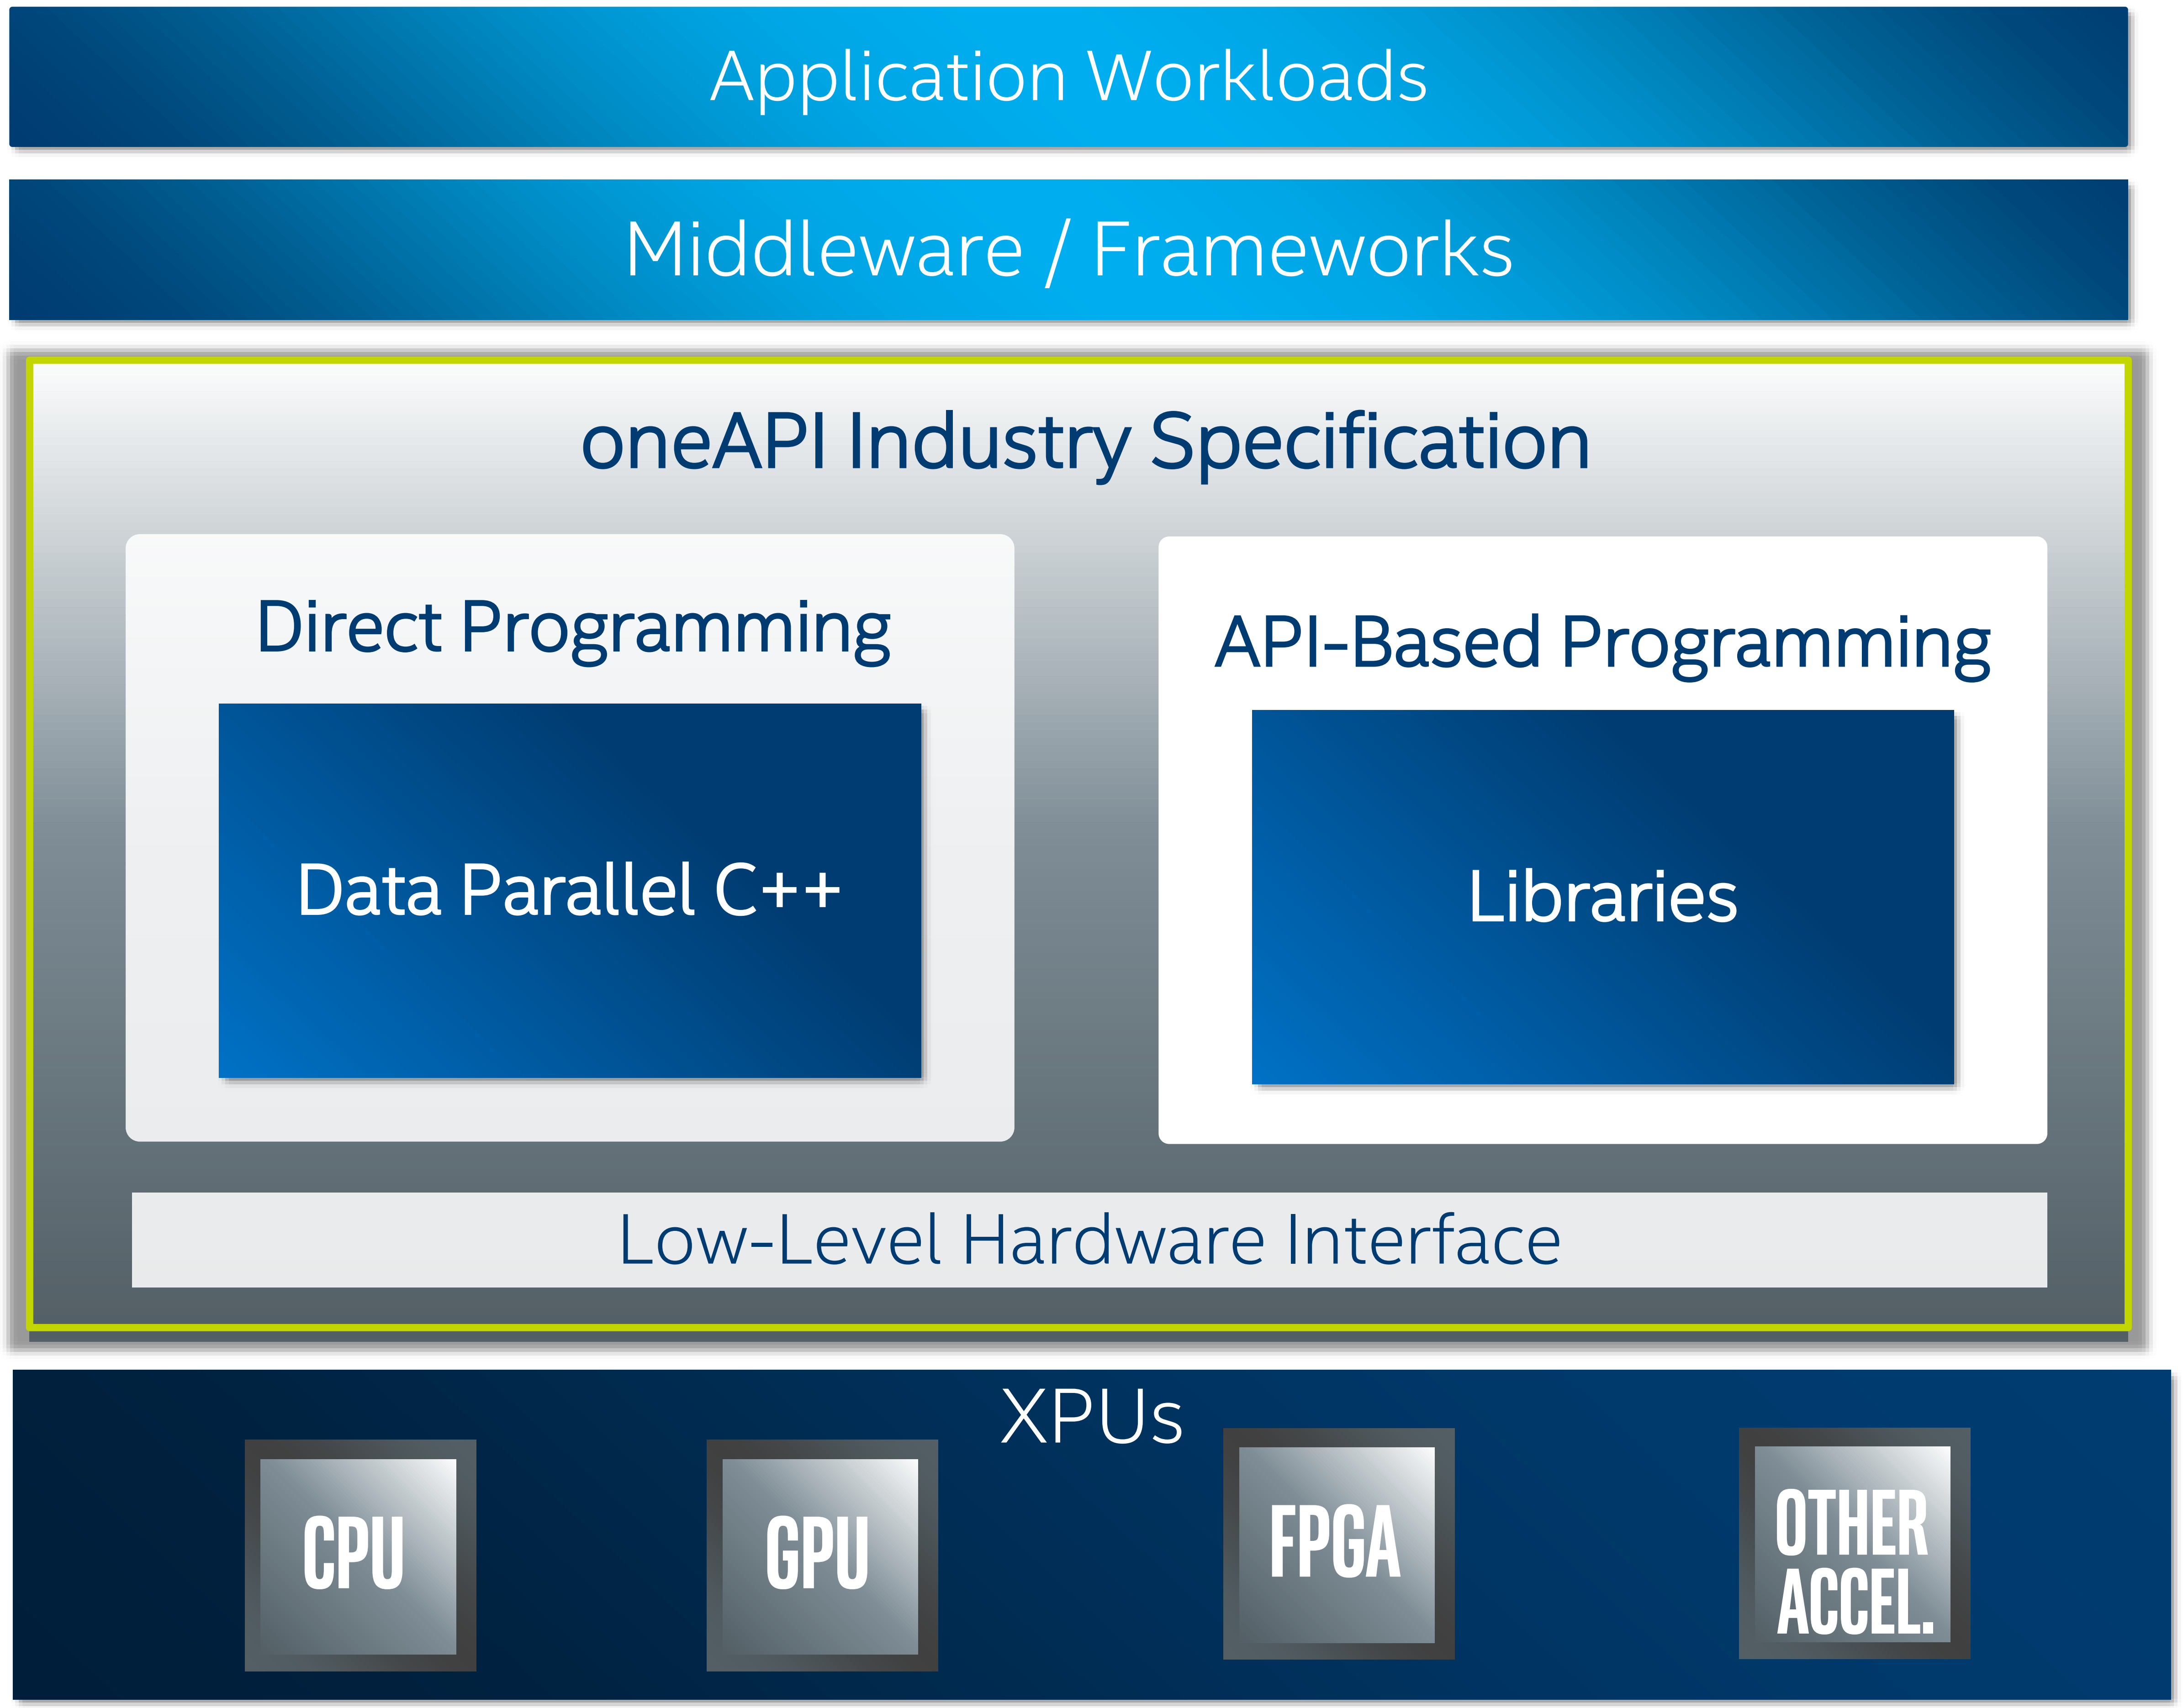

Visit oneapi.com for more details

Intel® oneAPI Programming Guide(Beta)

- https://software.intel.com/sites/default/files/oneAPIProgrammingGuide_5.pdf

Intel® oneAPI Programming Guide (Beta)

To support these needs, the code can be compiled into two main forms – a **SPIR-V** intermediate representation, which is the device-agnostic form, and target specific executable code.

- The SPIR-V representation enables online compilation, which is compilation to device code at execution time.
- Compilation to target specific executable code before executing the application is termed offline compilation. The Intel oneAPI DPC++ Compiler creates multiple code forms for each device.

These two forms are bundled together in an application binary known as a **fat binary**. The fat binary is employed by the oneAPI runtime when the application executed and form needed for the target device is determined at runtime. Device code for multiple devices co-exists in a single fat binary.

The following figure illustrates the compilation process from source code to fat binary.

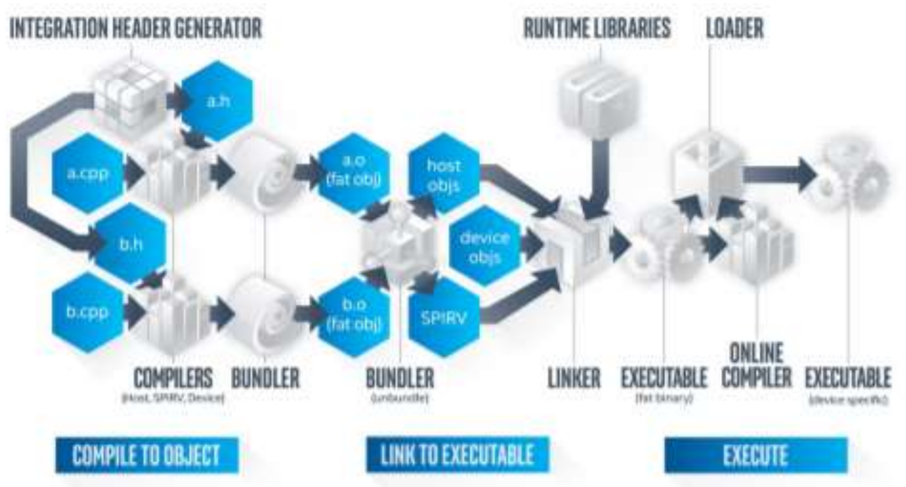

The tools participating in the compilation are represented by rectangles in the figure and are described as follows:

- **Driver** – The executable name is dpcpp on Linux* and dpcpp-cl on Windows*. Invoke this file on the command line to start compilation. It orchestrates the compilation process and invokes other tools as needed.
- **Host compiler** – the DPC++ host compiler. It is possible to employ a third-party host compiler by invoking the tools manually.
- **Device compiler** – a DPC++ device compiler (there can be more than one), SPIR-V compiler by default. The output files have the extension *.bc. SPIR-V is produced after the device code linkage step.
- **SPIRV compiler - llvm-spirv** – standard LLVM tool distributed with the Intel® oneAPI Base Toolkit that converts LLVM IR bitcode to SPIR-V.
- **Bundler** – Marshalls host and device object code into a single file called a **fat object**.


```
48      auto b_in =
  b_device.get_access<sycl::access::mode::read>(cgh);
49
50      cgh.parallel_for<class ex1>(a_size,[=](sycl::id<1> idx) {
51        c_res[idx] = a_in[idx] + b_in[idx];
52      });
53
54    });
55
56    }
57  }
```

A DPC++ program has the **single source** property, which means the **host code** and the **device code** can be placed in the same file so that the compiler treats them as the same compilation unit. This can potentially result in performance optimizations across the boundary between host and device code. The single source property differs from a programming model like OpenCL software technology where the host code and device code are typically in different files, and the host and device compiler are different entities, which means no optimization can occur between the host and device code boundary. Therefore, when scrutinizing a DPC++ program, the first step is to understand the delineation between host code and device code. To be more specific, DPC++ programs are delineated into different scopes similar to programming language scope, which is typically expressed via { and } in many languages.

The three types of scope in a DPC++ program include:

- **Application scope** – Code that executes on the host
- **Command group scope** – Code that acts as the interface between the host and device
- **Kernel scope** – Code that executes on the device

In this example, command group scope comprises lines 45 through 54 and includes coordination and data passing operations required in the program to enact control and communication between the host and the device.

```
45    d_queue.submit([&](sycl::handler &cgh) {
46      auto c_res = c_device.get_access<sycl::access::mode::write>(cgh);
47      auto a_in = a_device.get_access<sycl::access::mode::read>(cgh);
48      auto b_in = b_device.get_access<sycl::access::mode::read>(cgh);
49
50      cgh.parallel_for<class ex1>(a_size,[=](sycl::id<1> idx) {
51        c_res[idx] = a_in[idx] + b_in[idx];
52      });
53
54    });
```

Kernel scope, which is nested in the command group scope, comprises lines 50 to 52. Application scope consists of all the other lines not in command group or kernel scope. Syntactically, definitions are included from the top level include file; sycl.hpp and namespace declarations can be added for convenience.

インテル® oneAPIの始め方

- インテル® DevCloud、または、ローカル環境へのダウンロードの2通りの方法がございます。

ダウンロード＆インストール

https://software.intel.com/en-us/oneapi

インテル® DevCloud

https://software.intel.com/en-us/devcloud/oneapi

intel DevCloud

インテル® oneAPIツールキット(Beta)を使用して、かつ、インテルCPU、GPU、FPGAをまたがったワークロードを開発、テスト、実行するためのサンドボックス

- インテル® oneAPI ツールキット の使用
- Data Parallel C++の学習
- ワークロードの評価
- ヘテロジニアスアプリケーション構築
- プロジェクトのプロトタイピング

No downloads | No hardware acquisition | No installation | No set-up & configuration

Get up & running in seconds!

intel

We are in a data-centric world

All data must be

stored, processed, and analyzed

異なるデータ階層ごとの技術革新の必要性

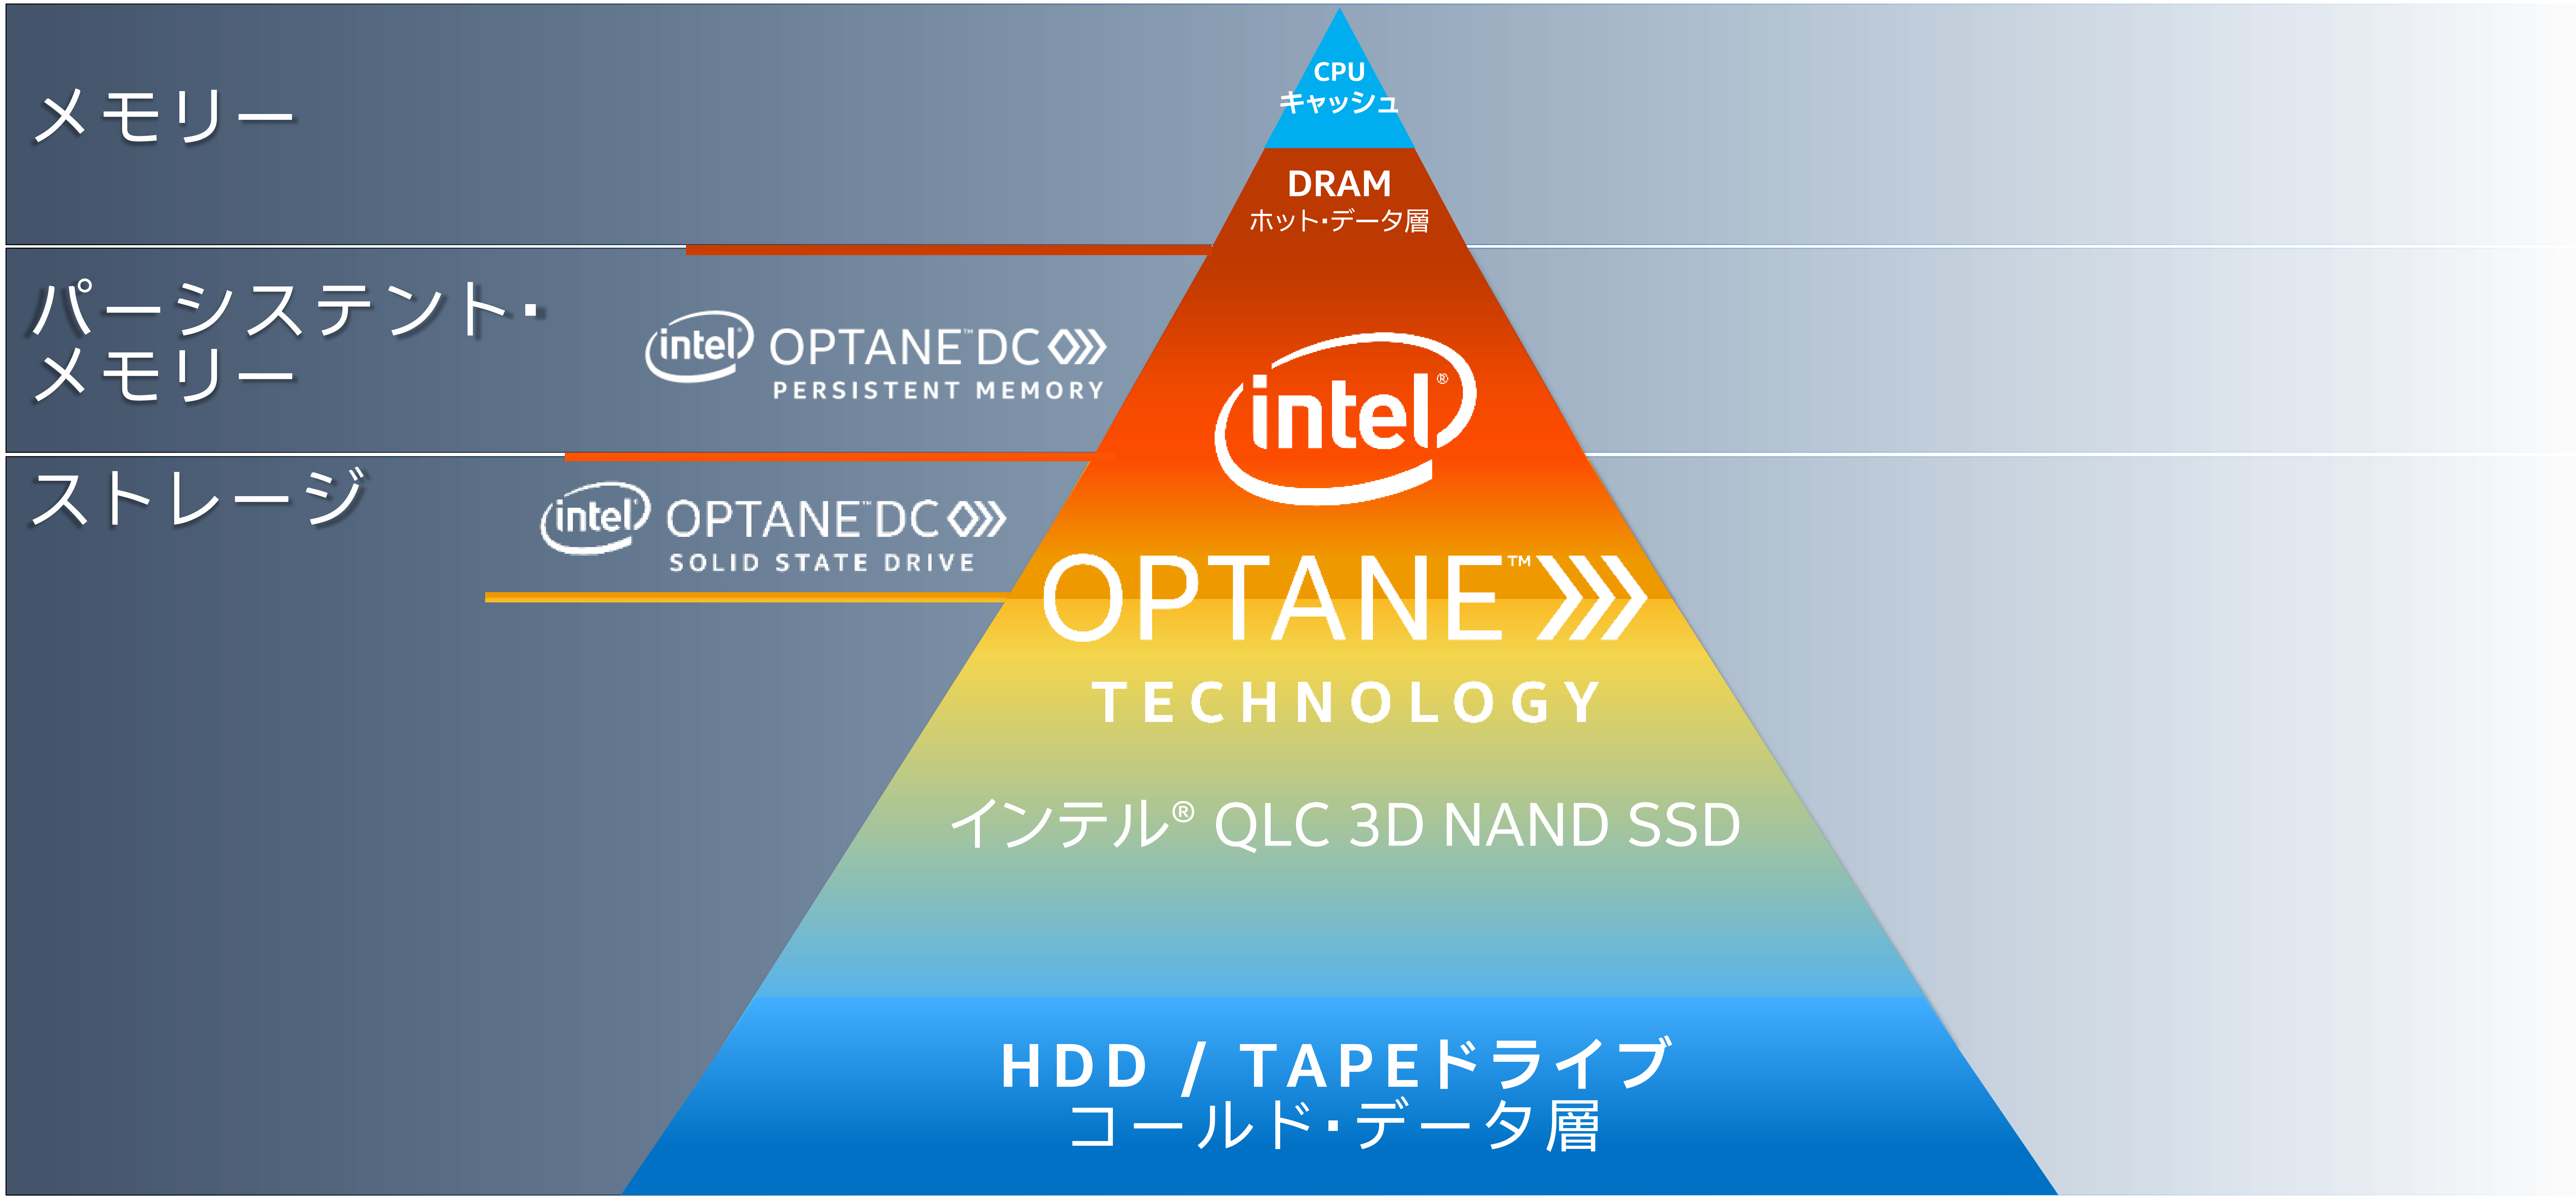

更なるNANDの階層化

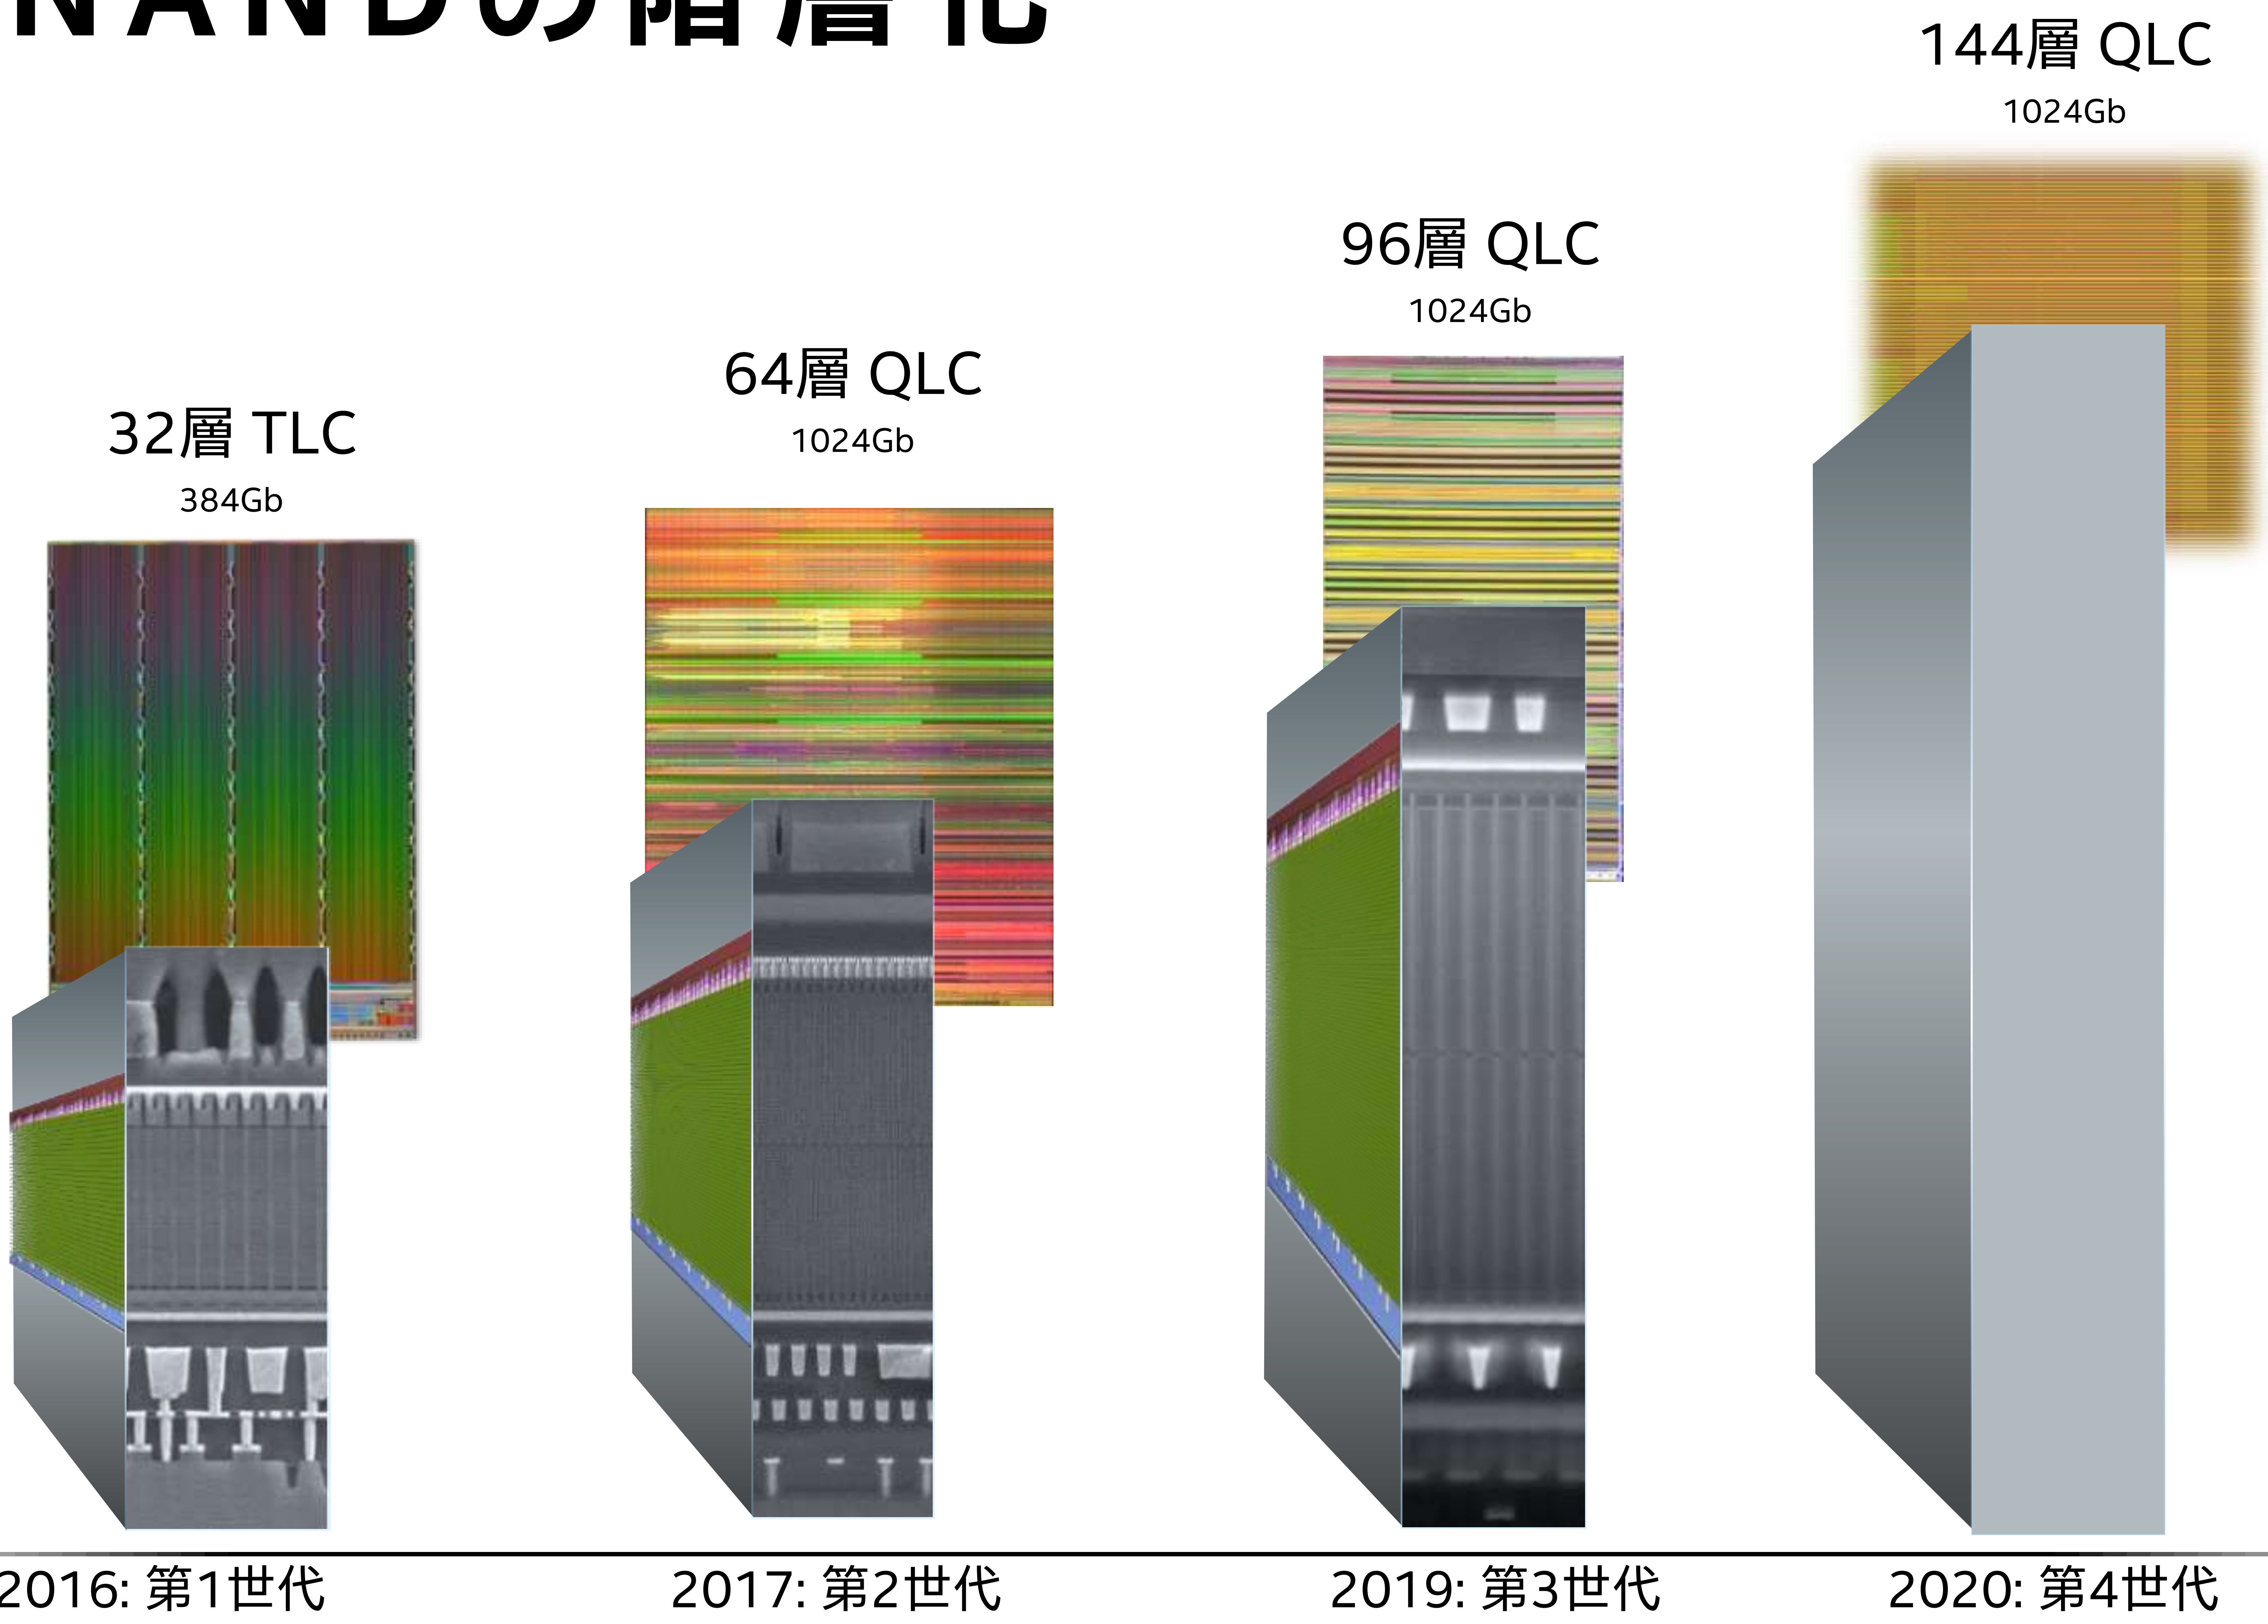

インテルのフローティング･ゲート･セルは 5 ビット / セルへと拡張

5 ビット / セル
32 states

4 ビット / セル
16 states

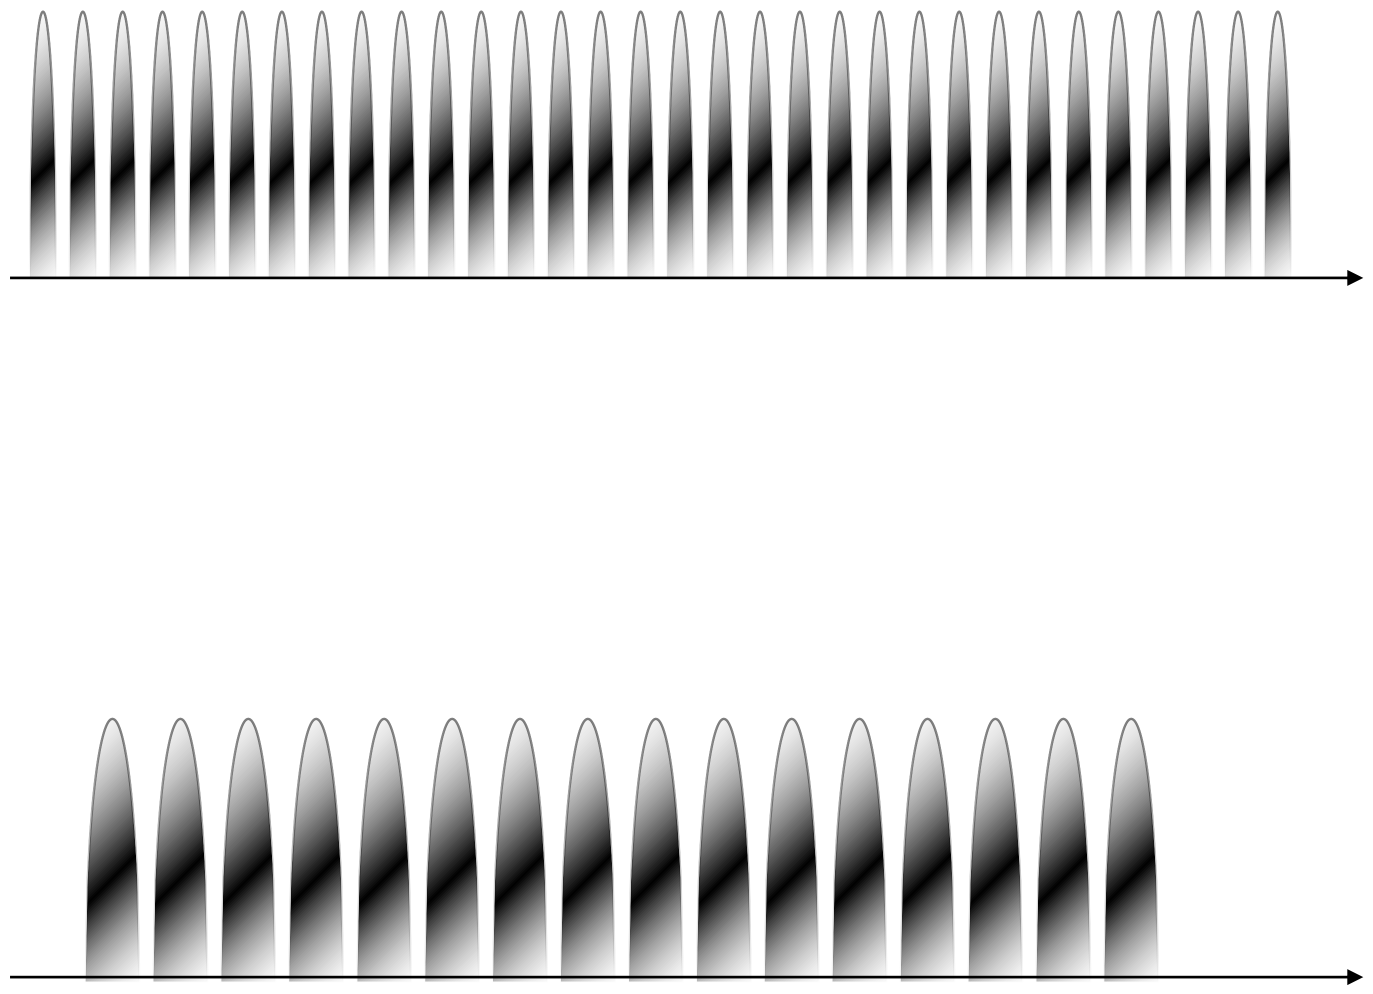

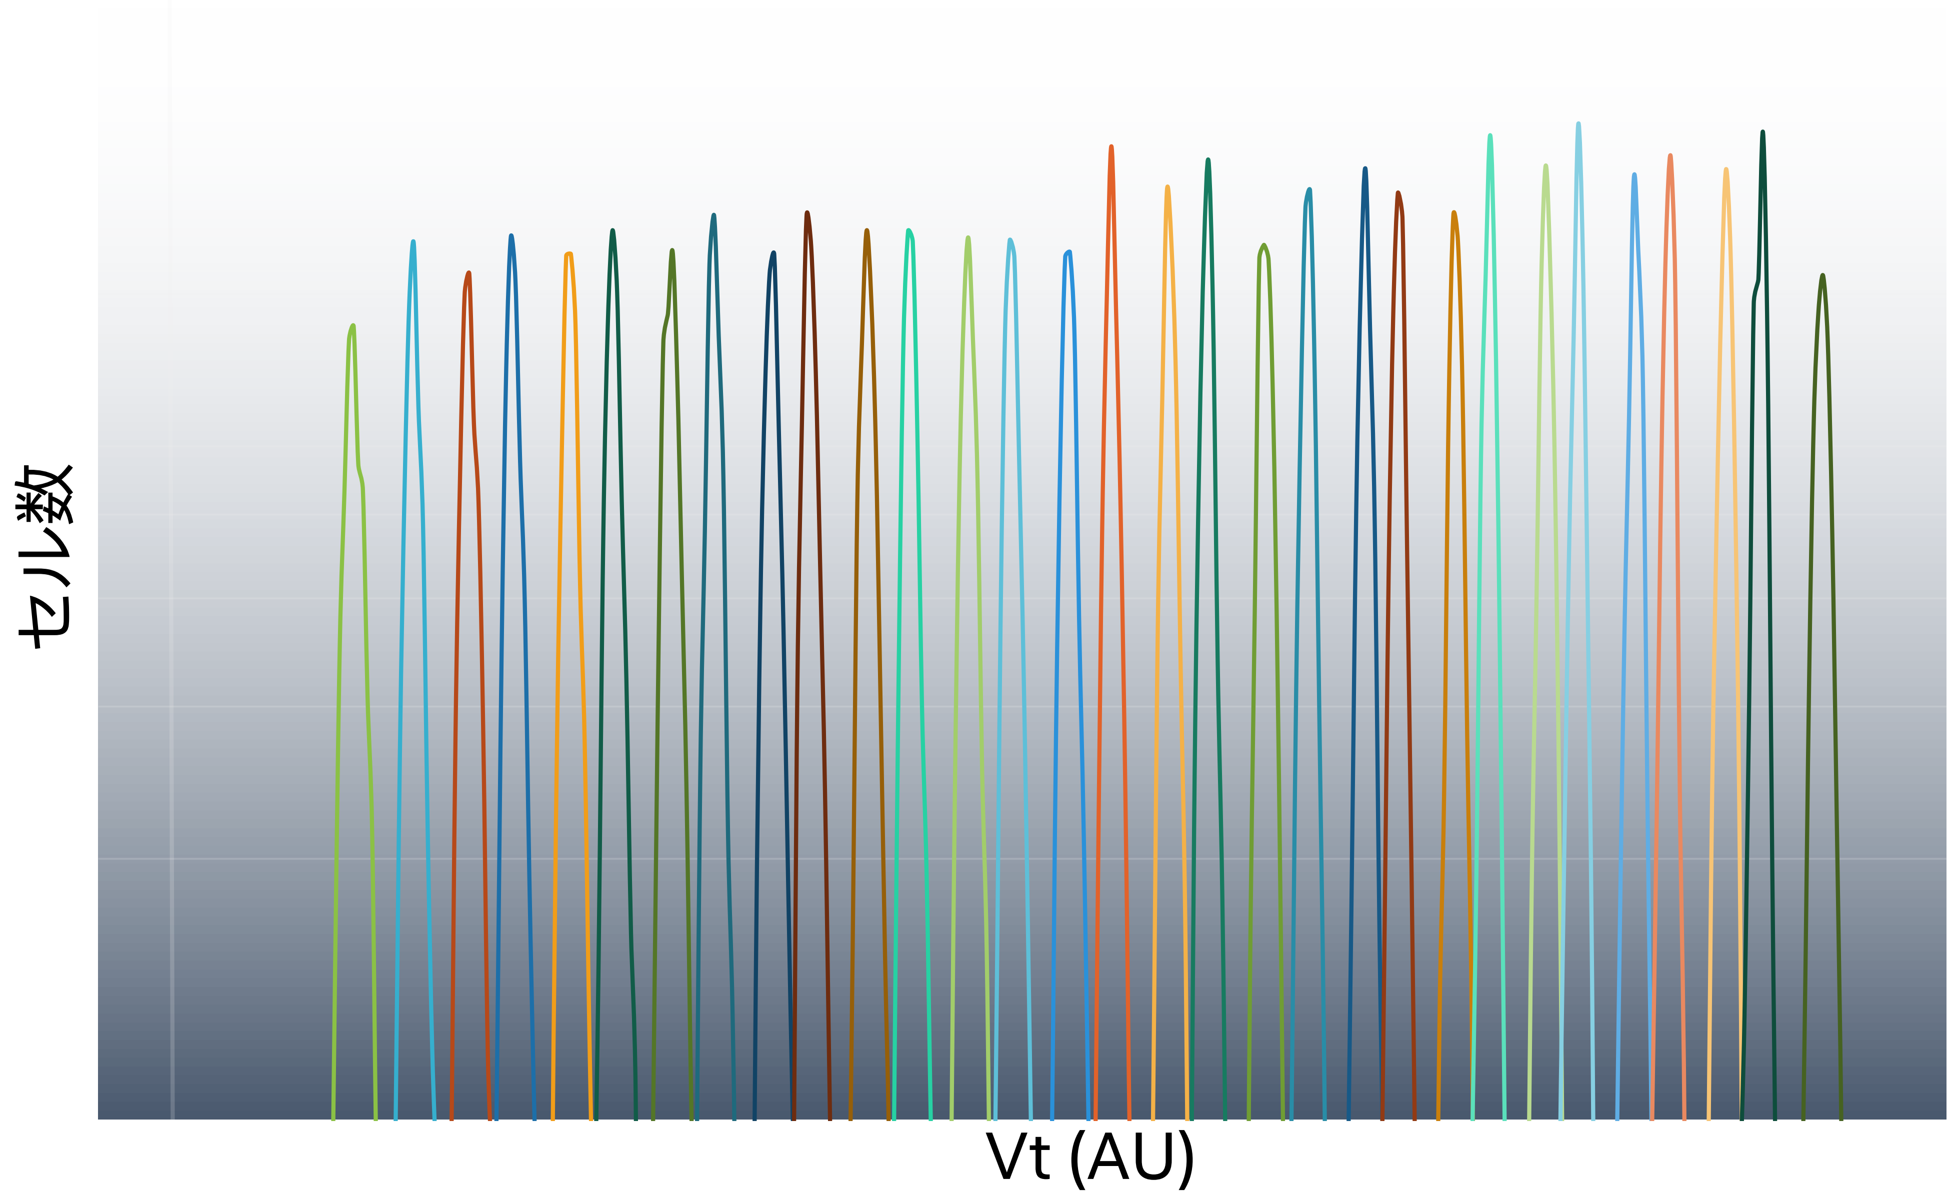

より高密度な実装へ:
ビット / セル→ ビット / ダイ → ビット / ウエハー → ビット / SSD → ビット / ラック

Highest Density Media Meets Highest Density Form Factor

SOLUTION RANGE

E1.L 18mm and 9.5mm
E1.S, 5.9mm

THERMAL EFFICIENCY

E1.L 2x more thermally efficient than U.2 15mm[2]
E1.S 3x more thermally efficient than U.2 7mm[3]

CAPACITY SCALING

E1.L up to 2.6x the capacity per 1 rack unit than U.2[1]
E1.S up to 2x the capacity per drive than M.2[4]

FUTURE READY

PCIe* 4.0 and 5.0 ready[5]

EDSFF specifications https://edsffspec.org/edsff-resources/
See footnotes and disclaimers in Appendix A.
For more complete information about performance and benchmark results, visit www.intel.com/benchmarks

What is Intel® Optane™ Technology?

TRANSISTOR-LESS DESIGN

Data is written at a bit level, so **each cell's state can be changed from a 0 or 1 independently** of other cells

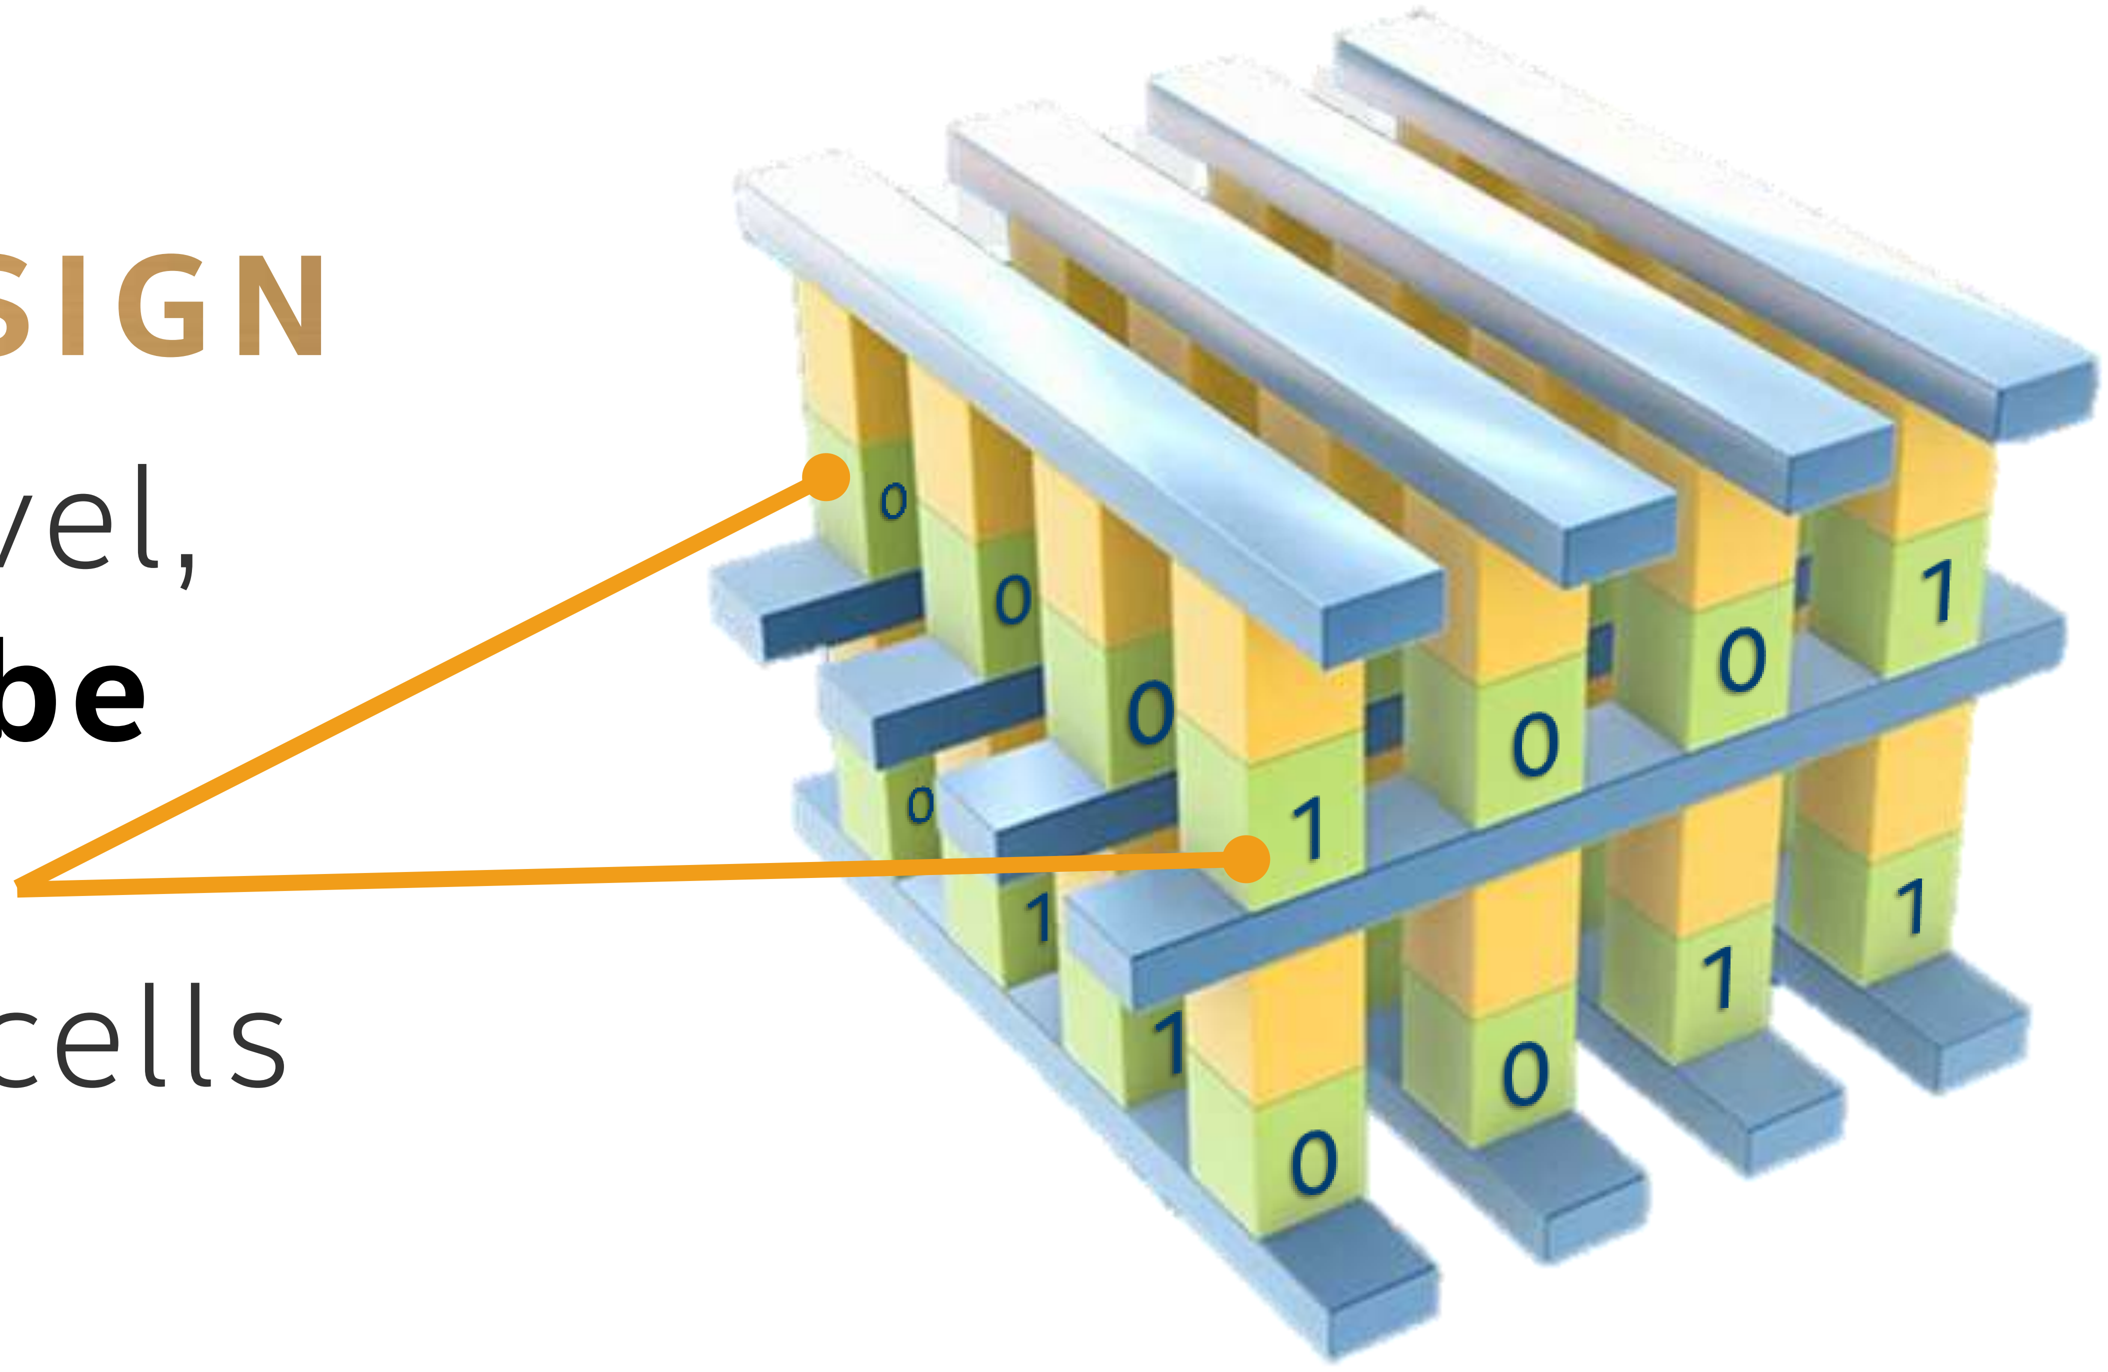

Intel® Optane™ Memory Media

Intel® Optane™ Technology design is fundamentally different from NAND

Intel® Optane™ DC Technology Products Coexist

Improving **Memory Capacity**

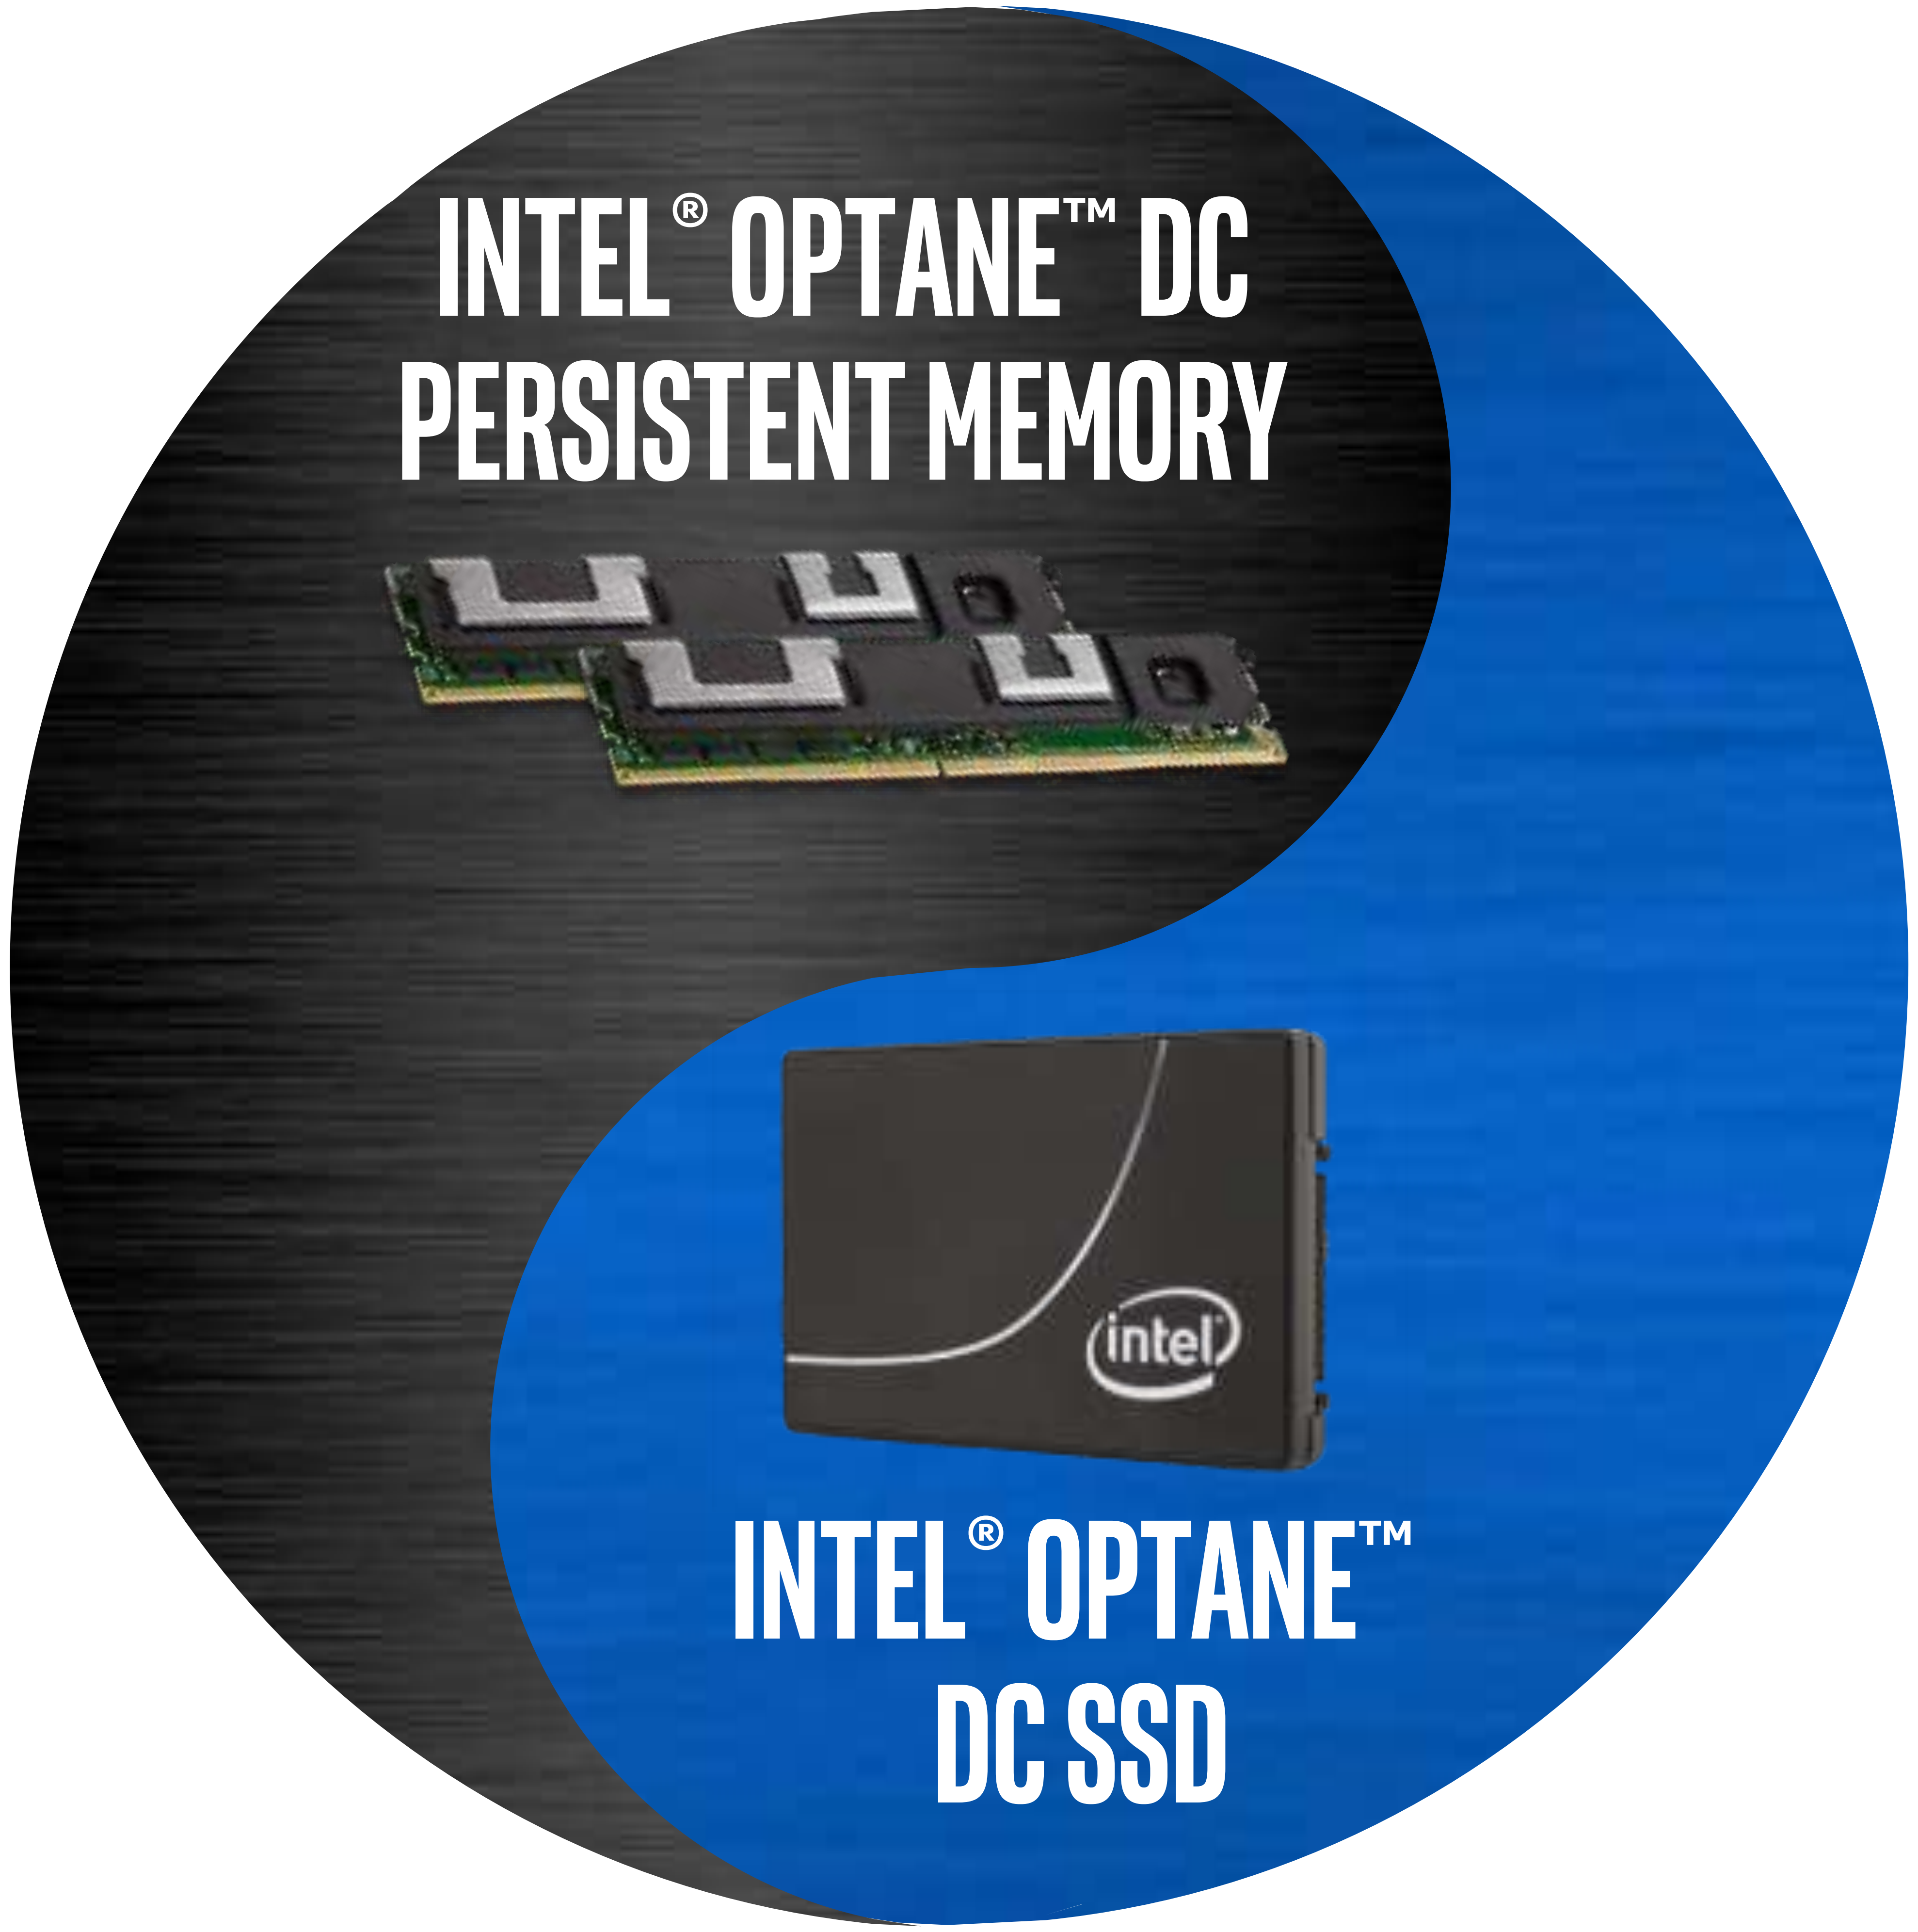

Improving **Working Storage**

More to be Gained by being on memory Bus

ランダムリード平均アイドルレイテンシー

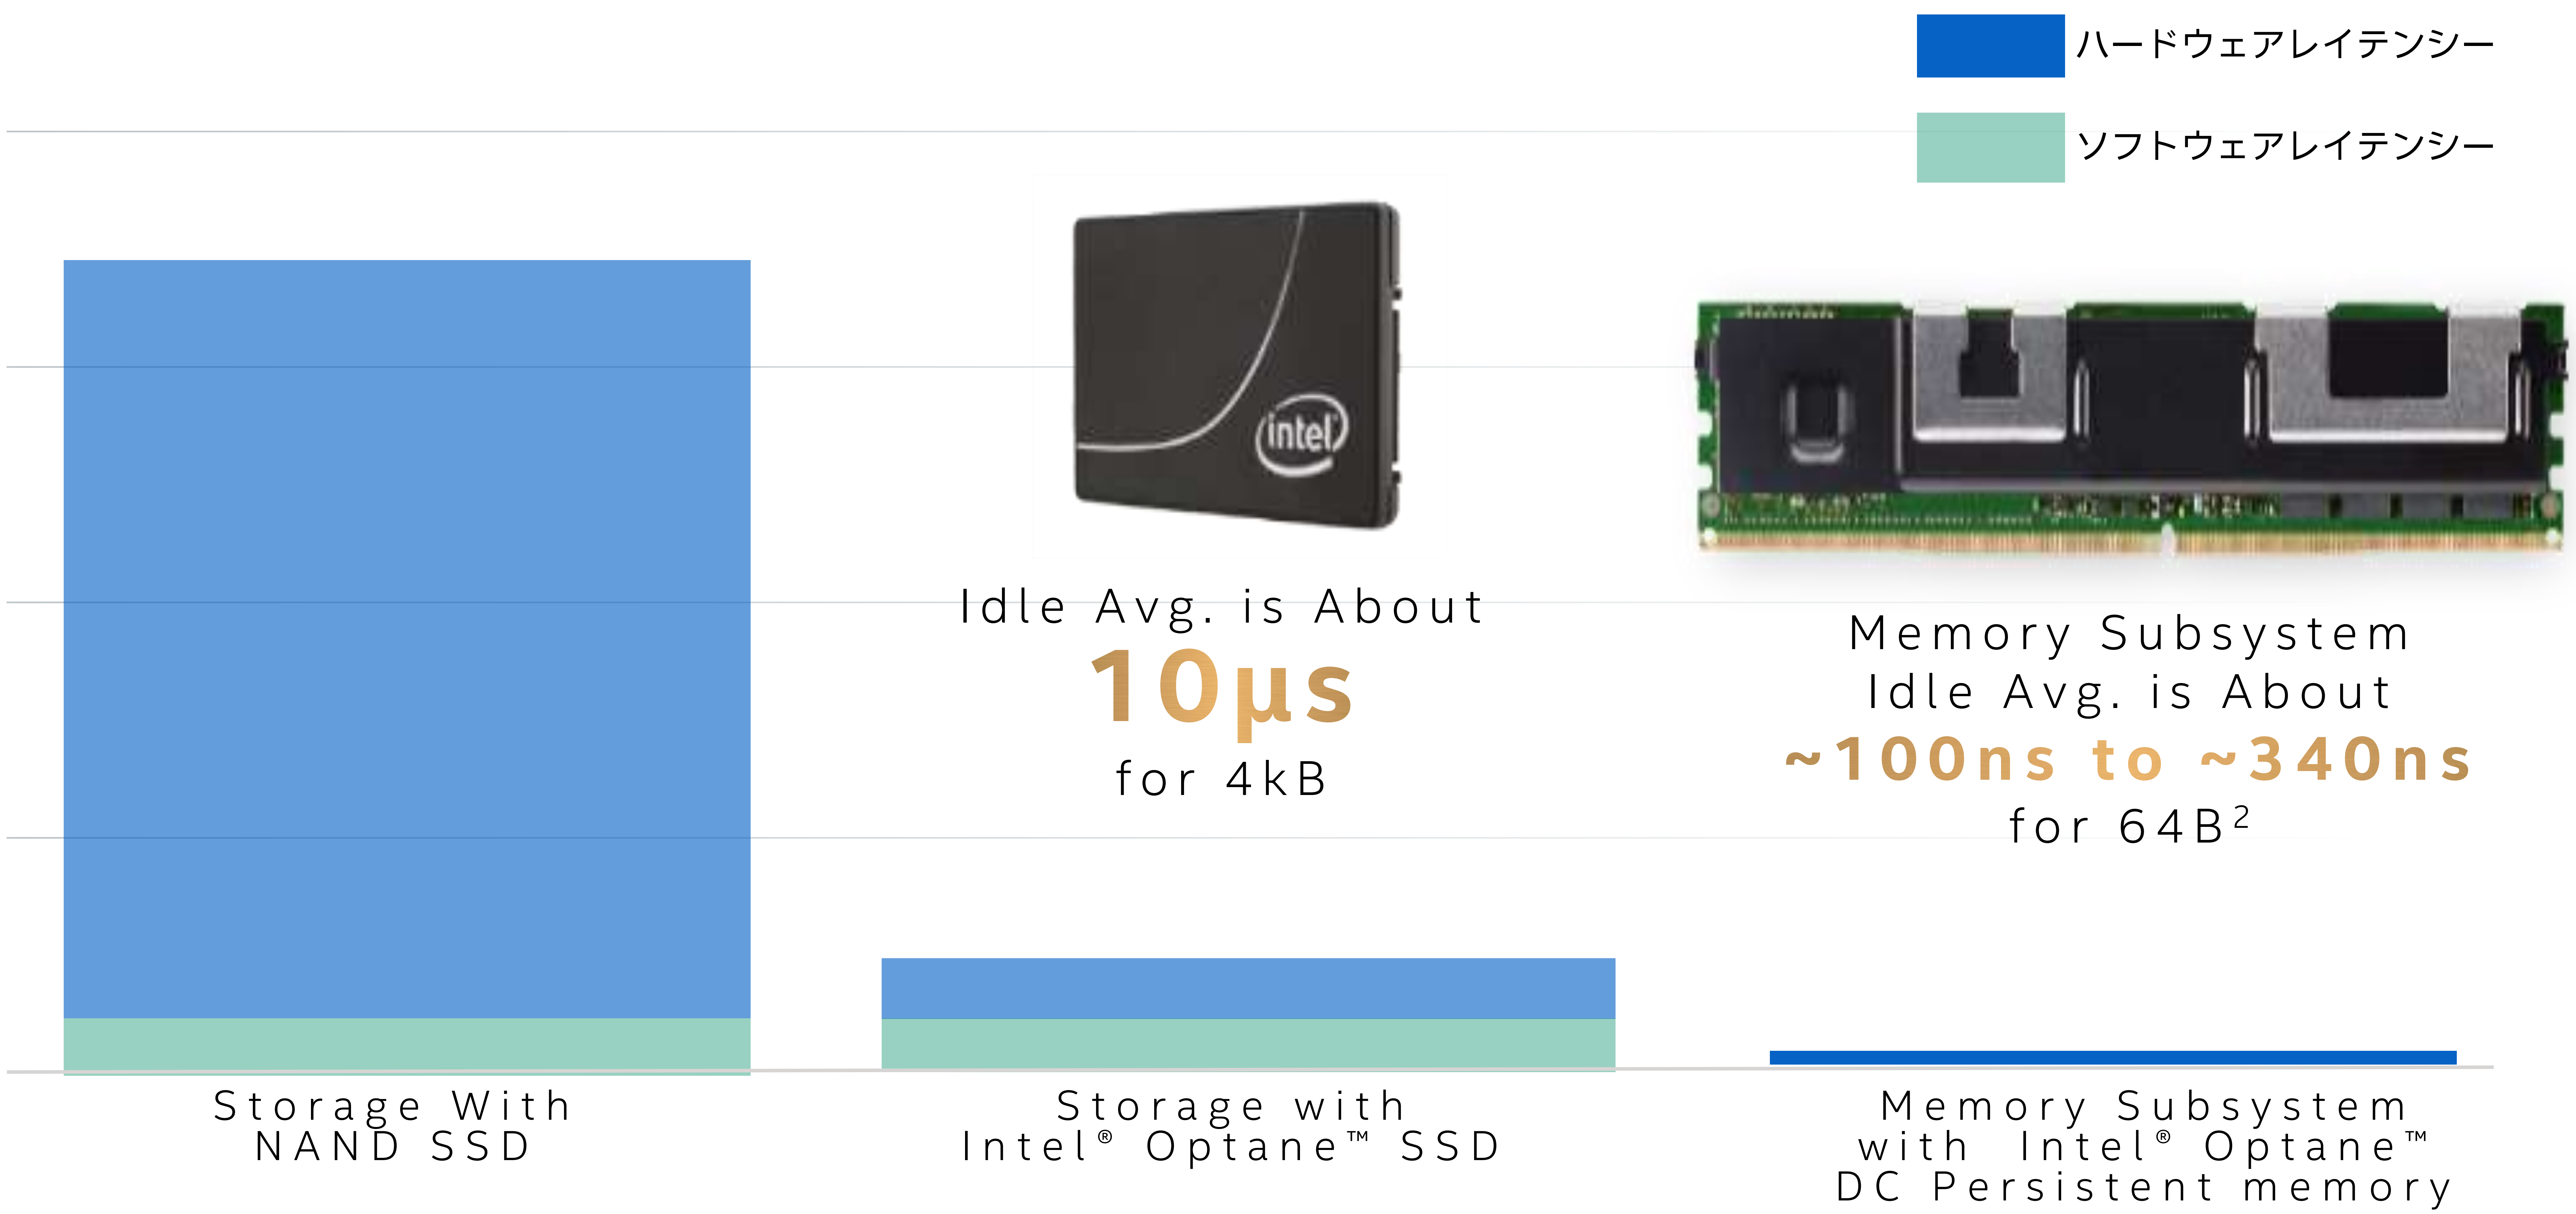

Aerospike Certification Test (ACT) RESULTS

Test read latency under 1 ms at high write pressure for key-value database usage

AEROSPIKE CERTIFICATION TEST FAILURE RATE @ 300K TPS

ALDERSTREAM VS CURRENT INTEL® SSDS (LOWER IS BETTER)

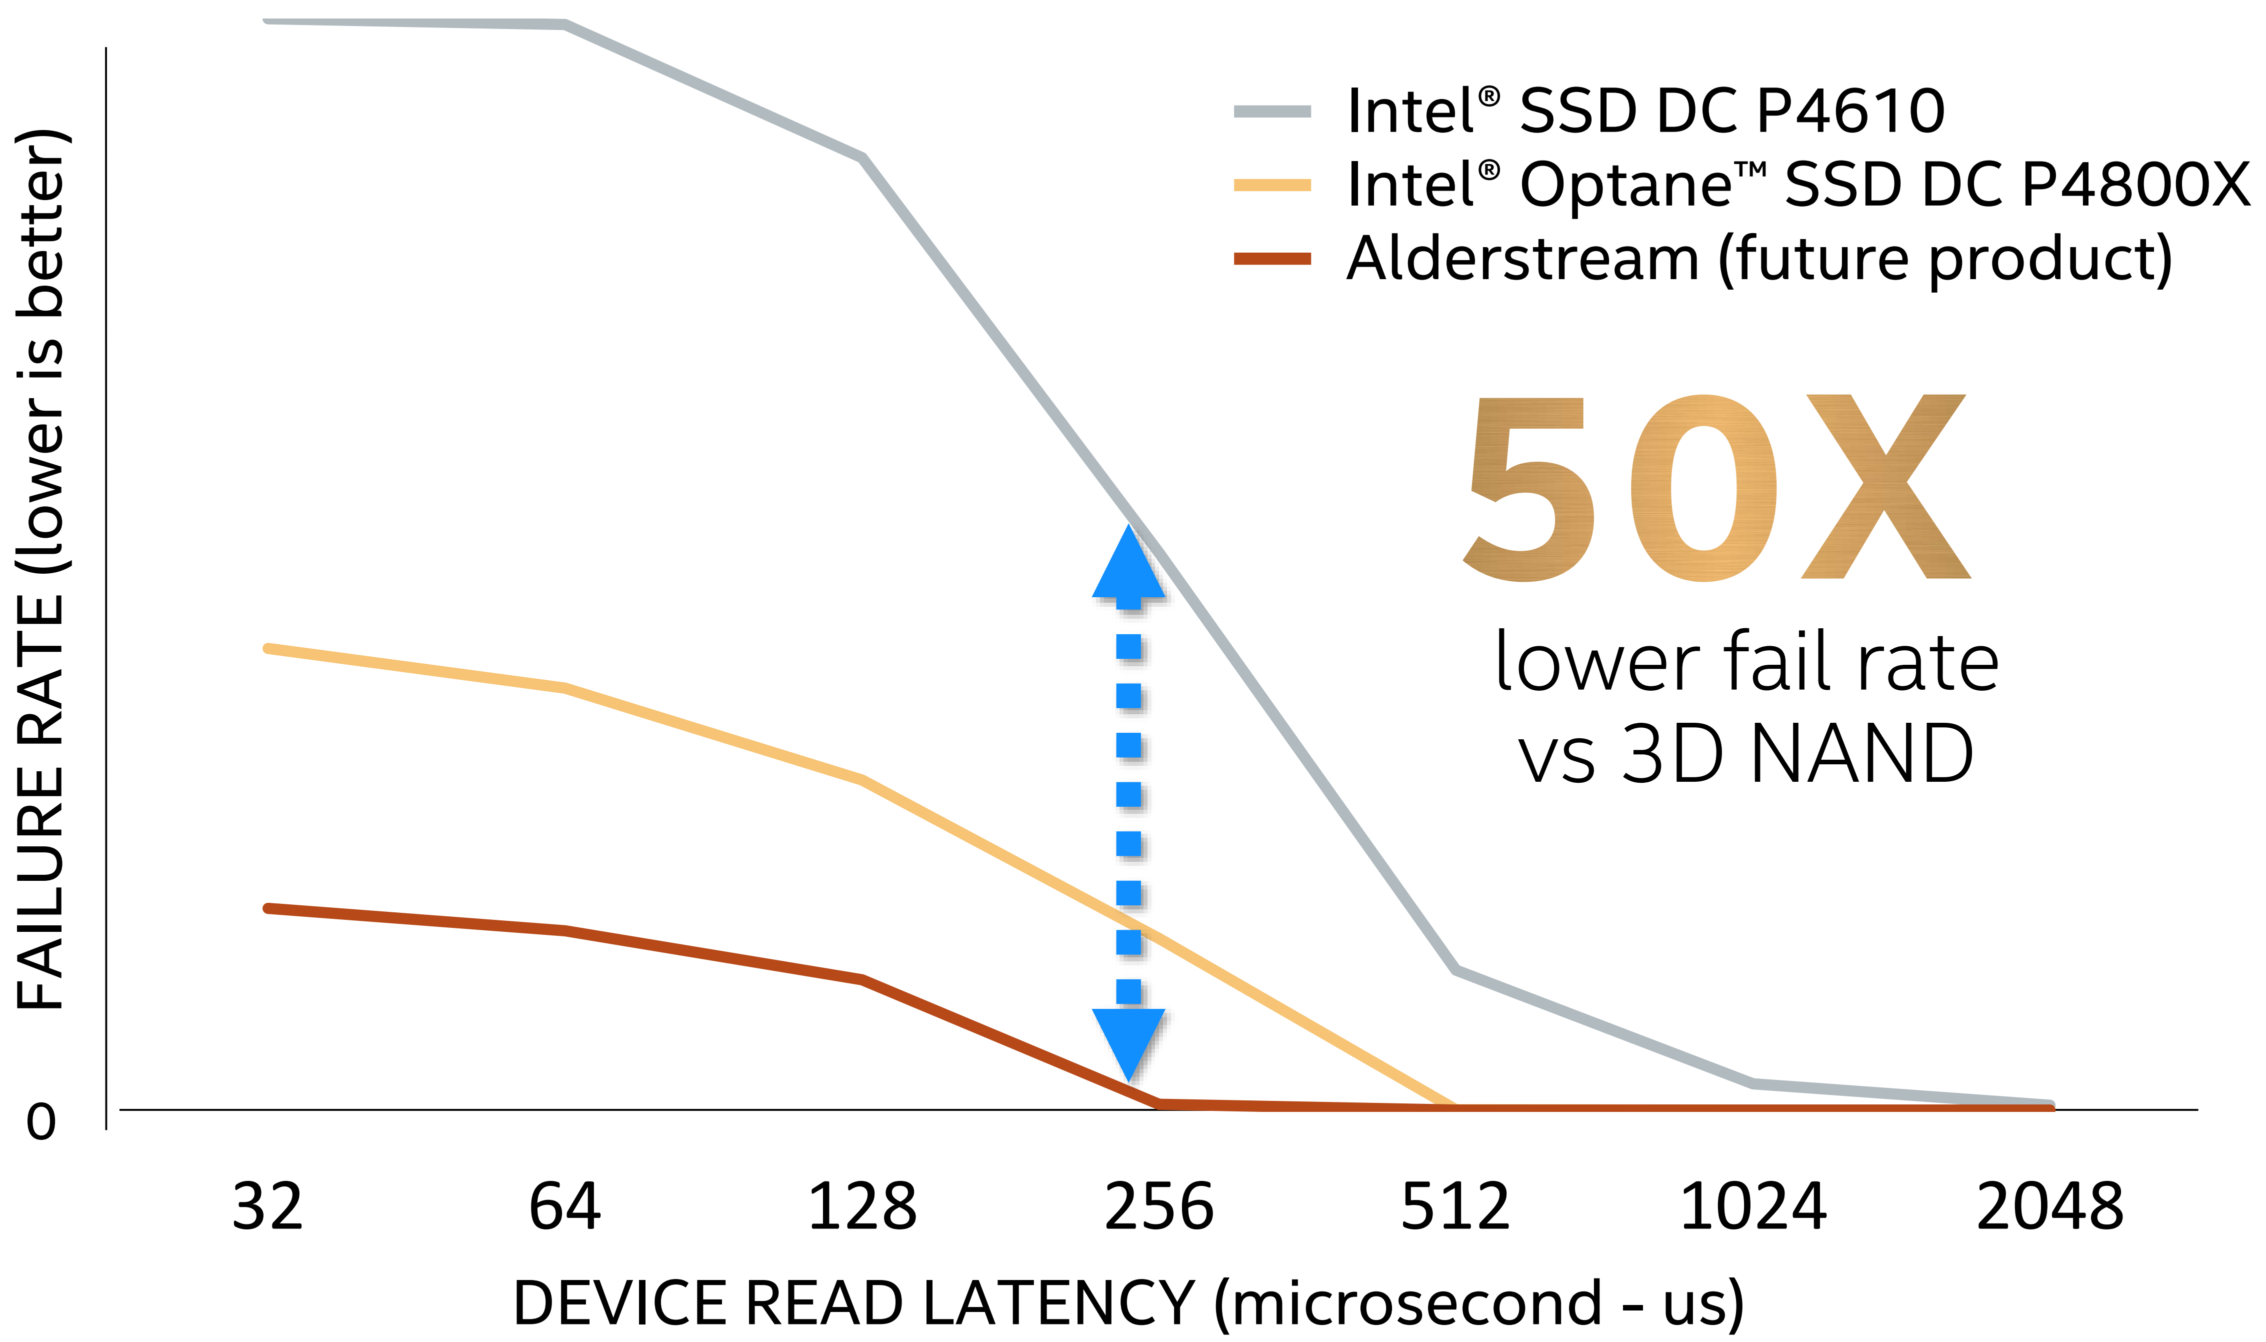

Maximum TPS at <5% ACT Failure Rate

(HIGHER IS BETTER)

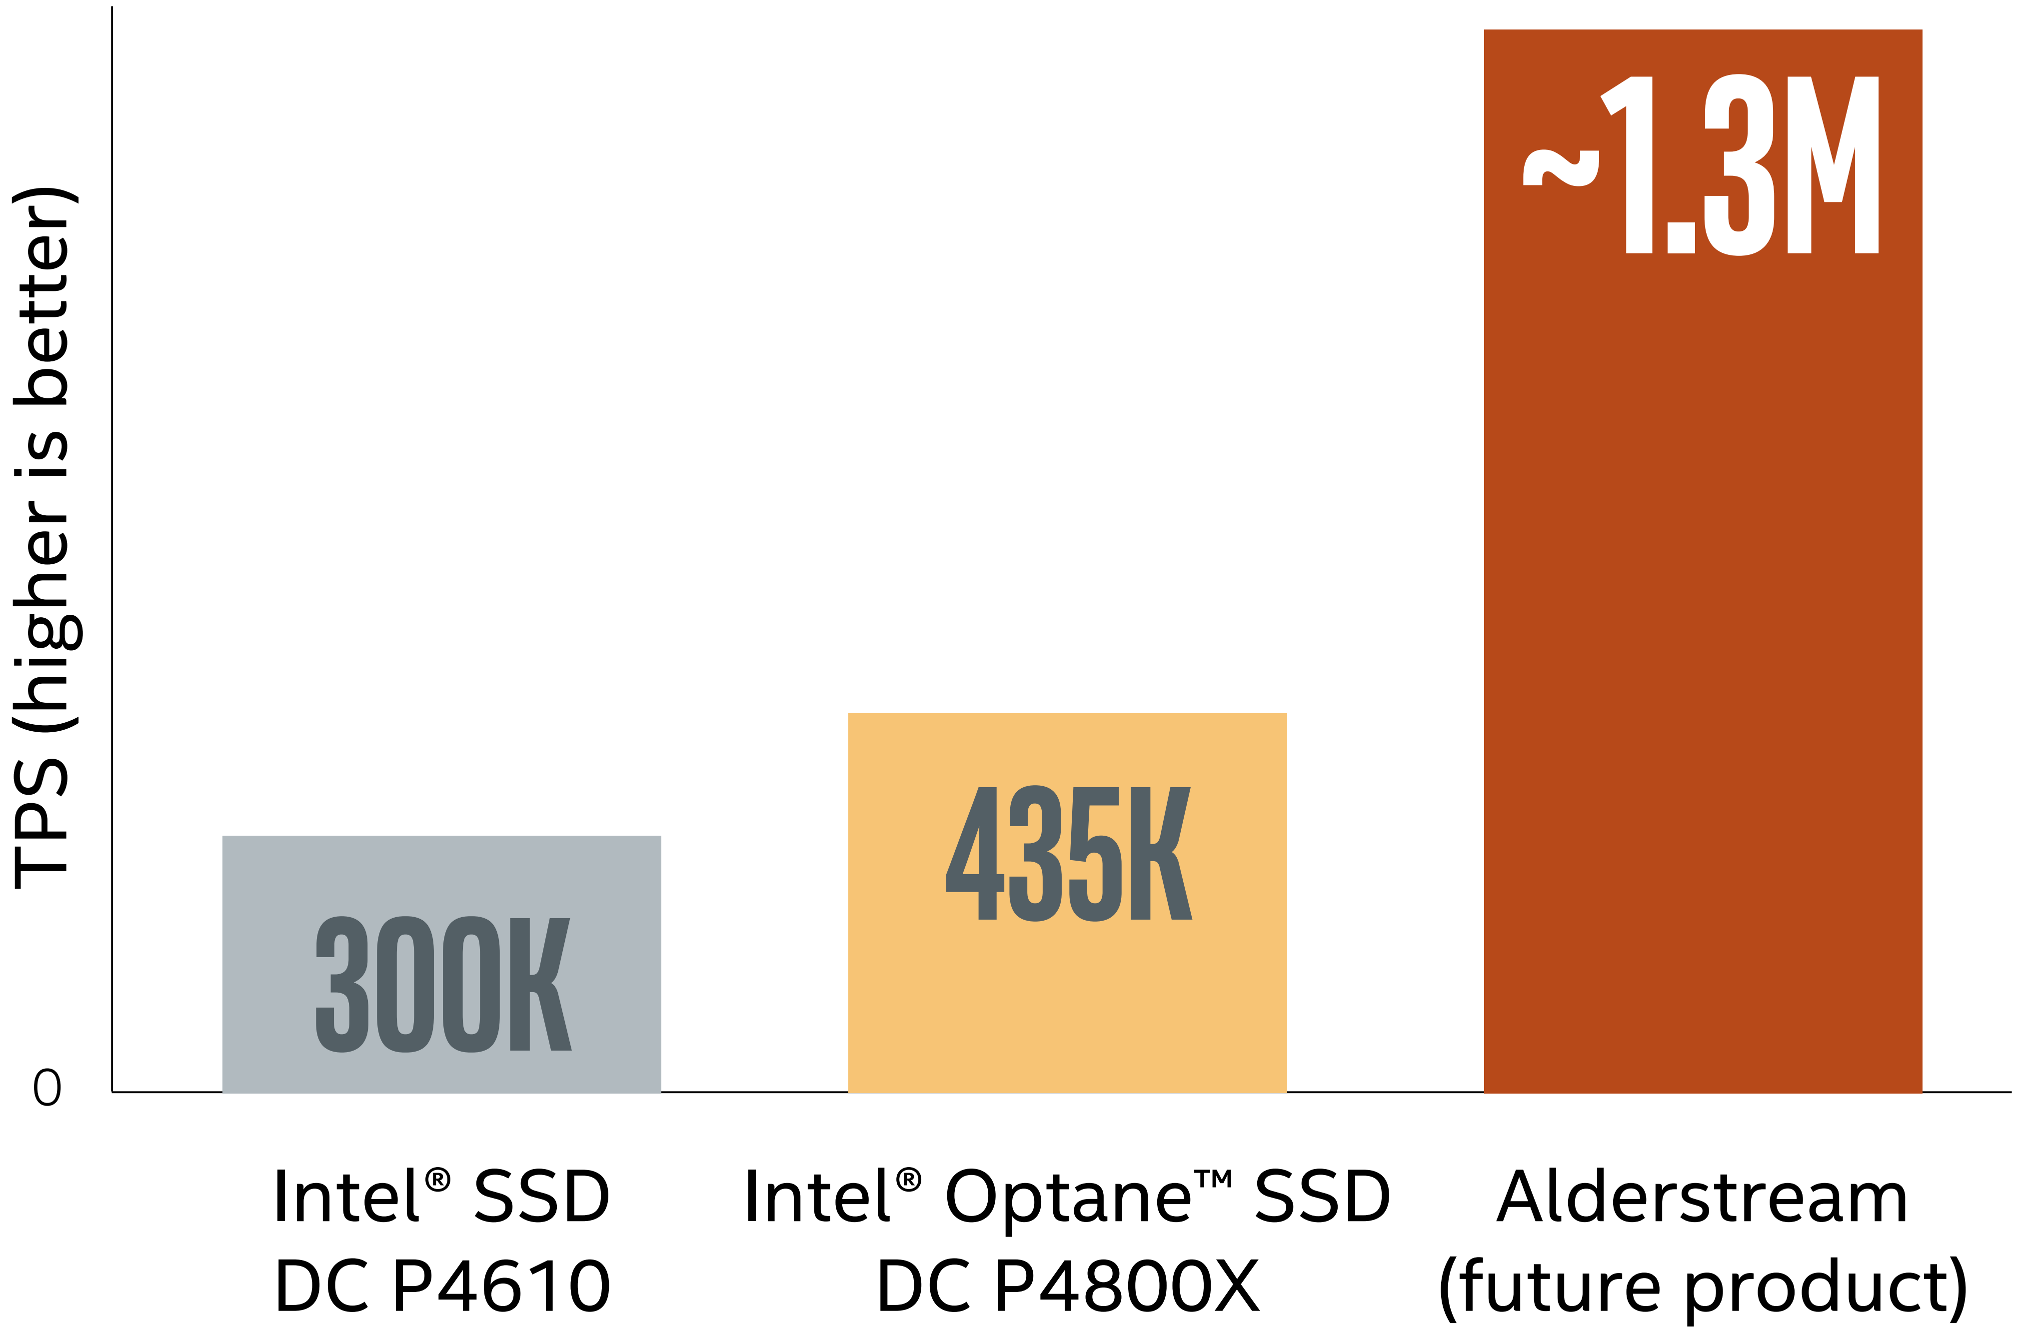

Bottlenecks:
The Nemesis of HPC Performance

HPC Clusters are only as fast as their slowest component

Challenge:	Breakthroughs:
Compute	Intel® Xeon® processors Intel® FPGAs
Network	Dual-ported 100Gb
Storage Media	Intel® 3D NAND SSDs Intel® Optane™ SSDs
POSIX block-based I/O	No good solution… until now

インテル® Optane™ DCパーシステント・メモリー

メモリーおよびストレージのパラダイムの変更

HPC ストレージモデル

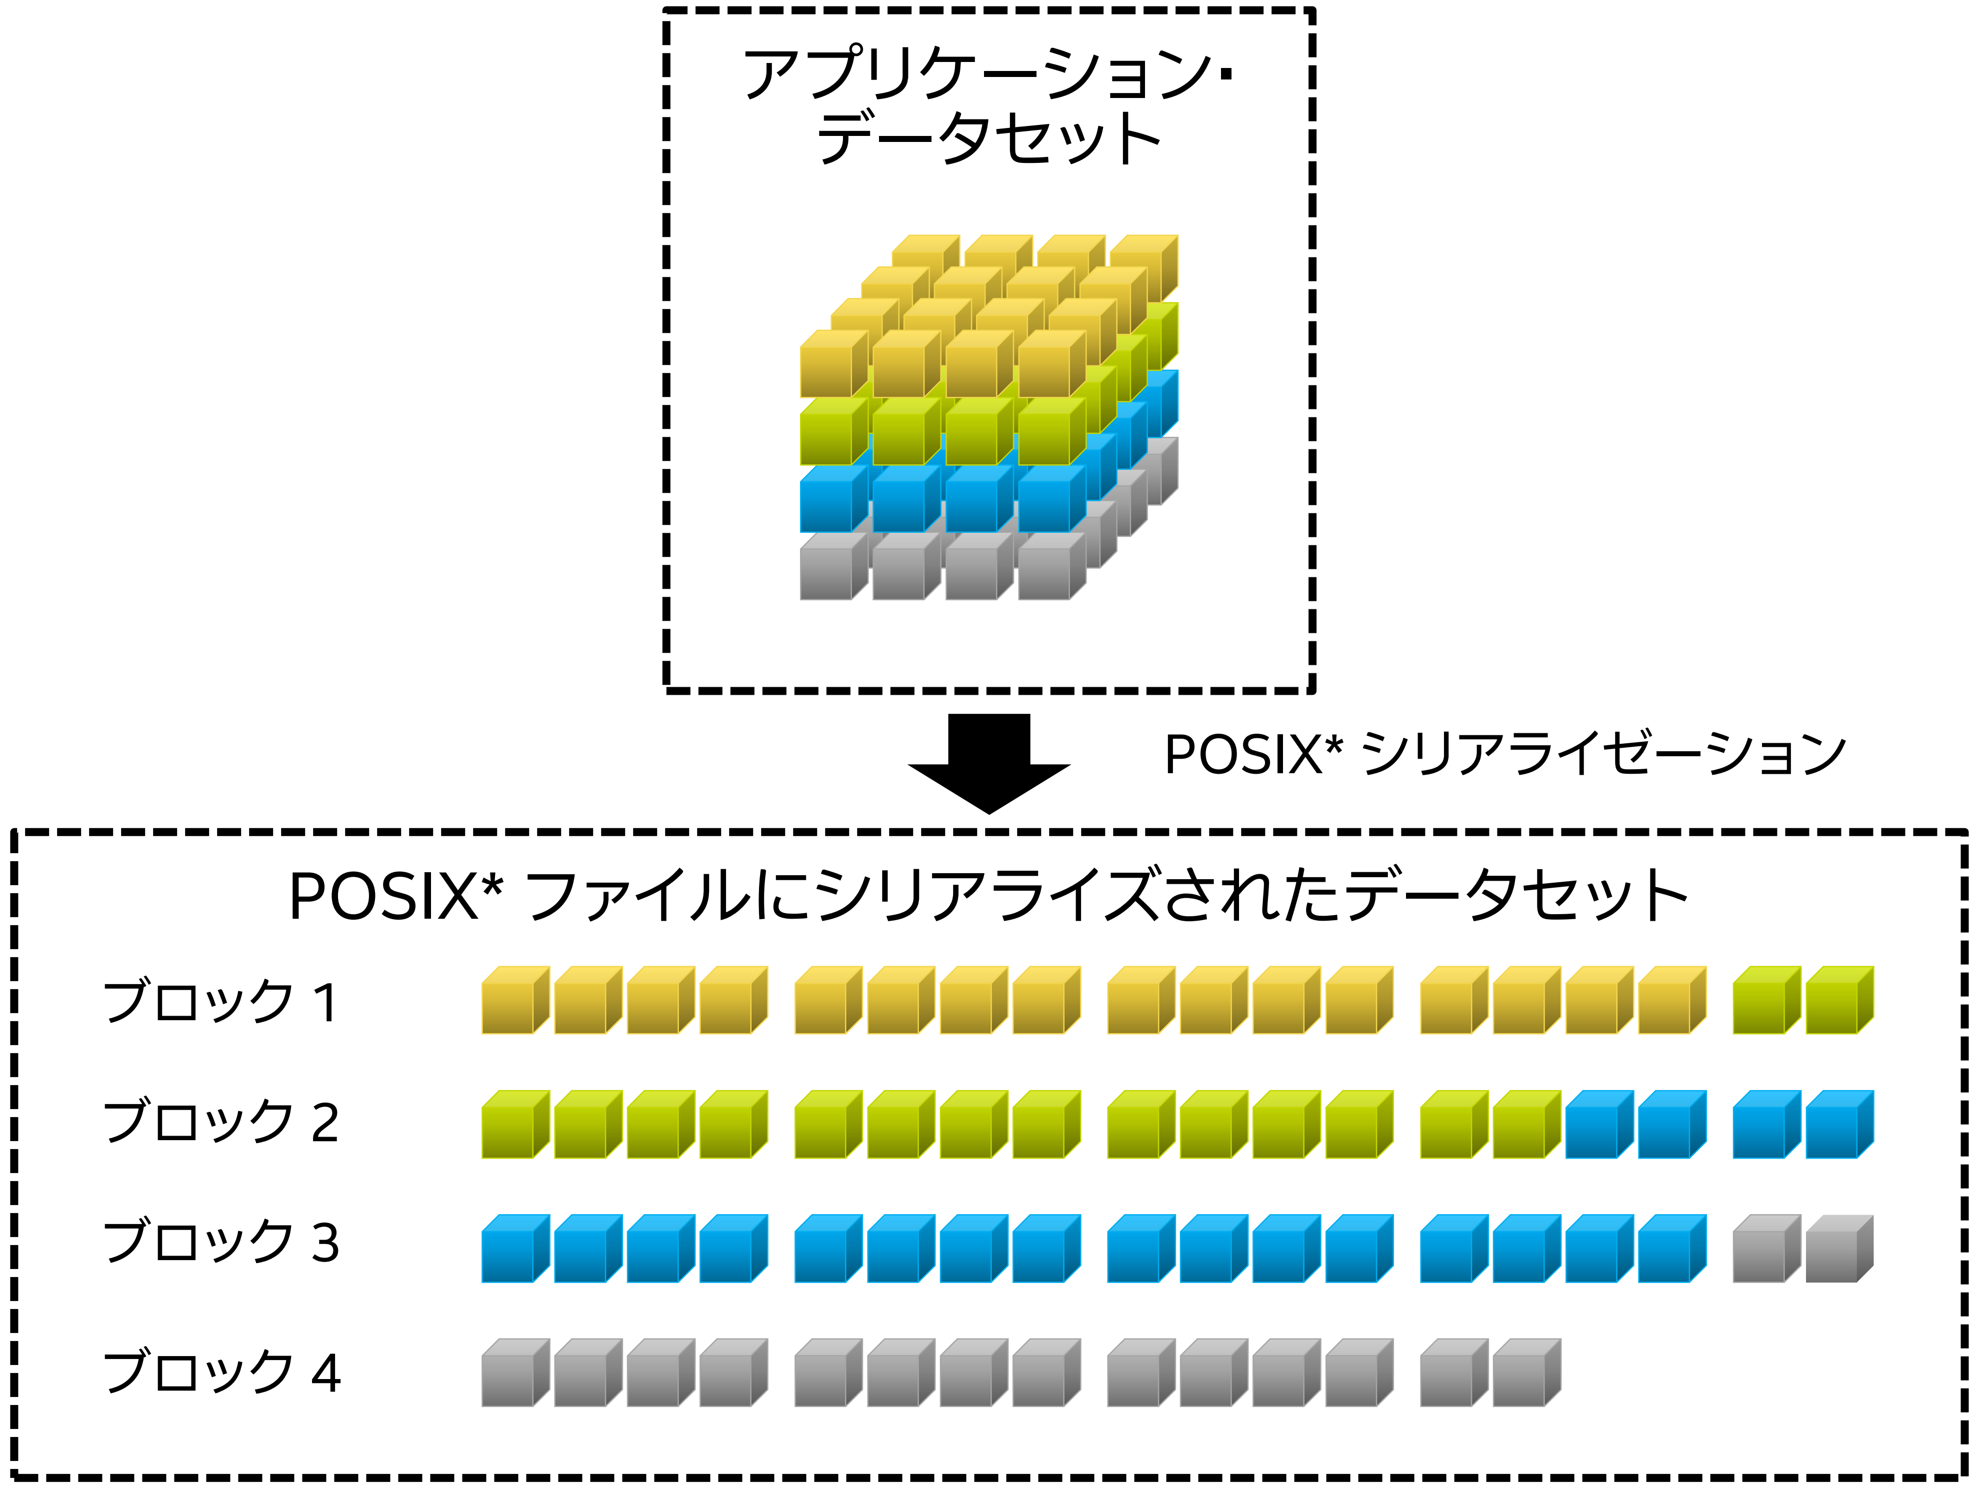

AI ストレージモデル

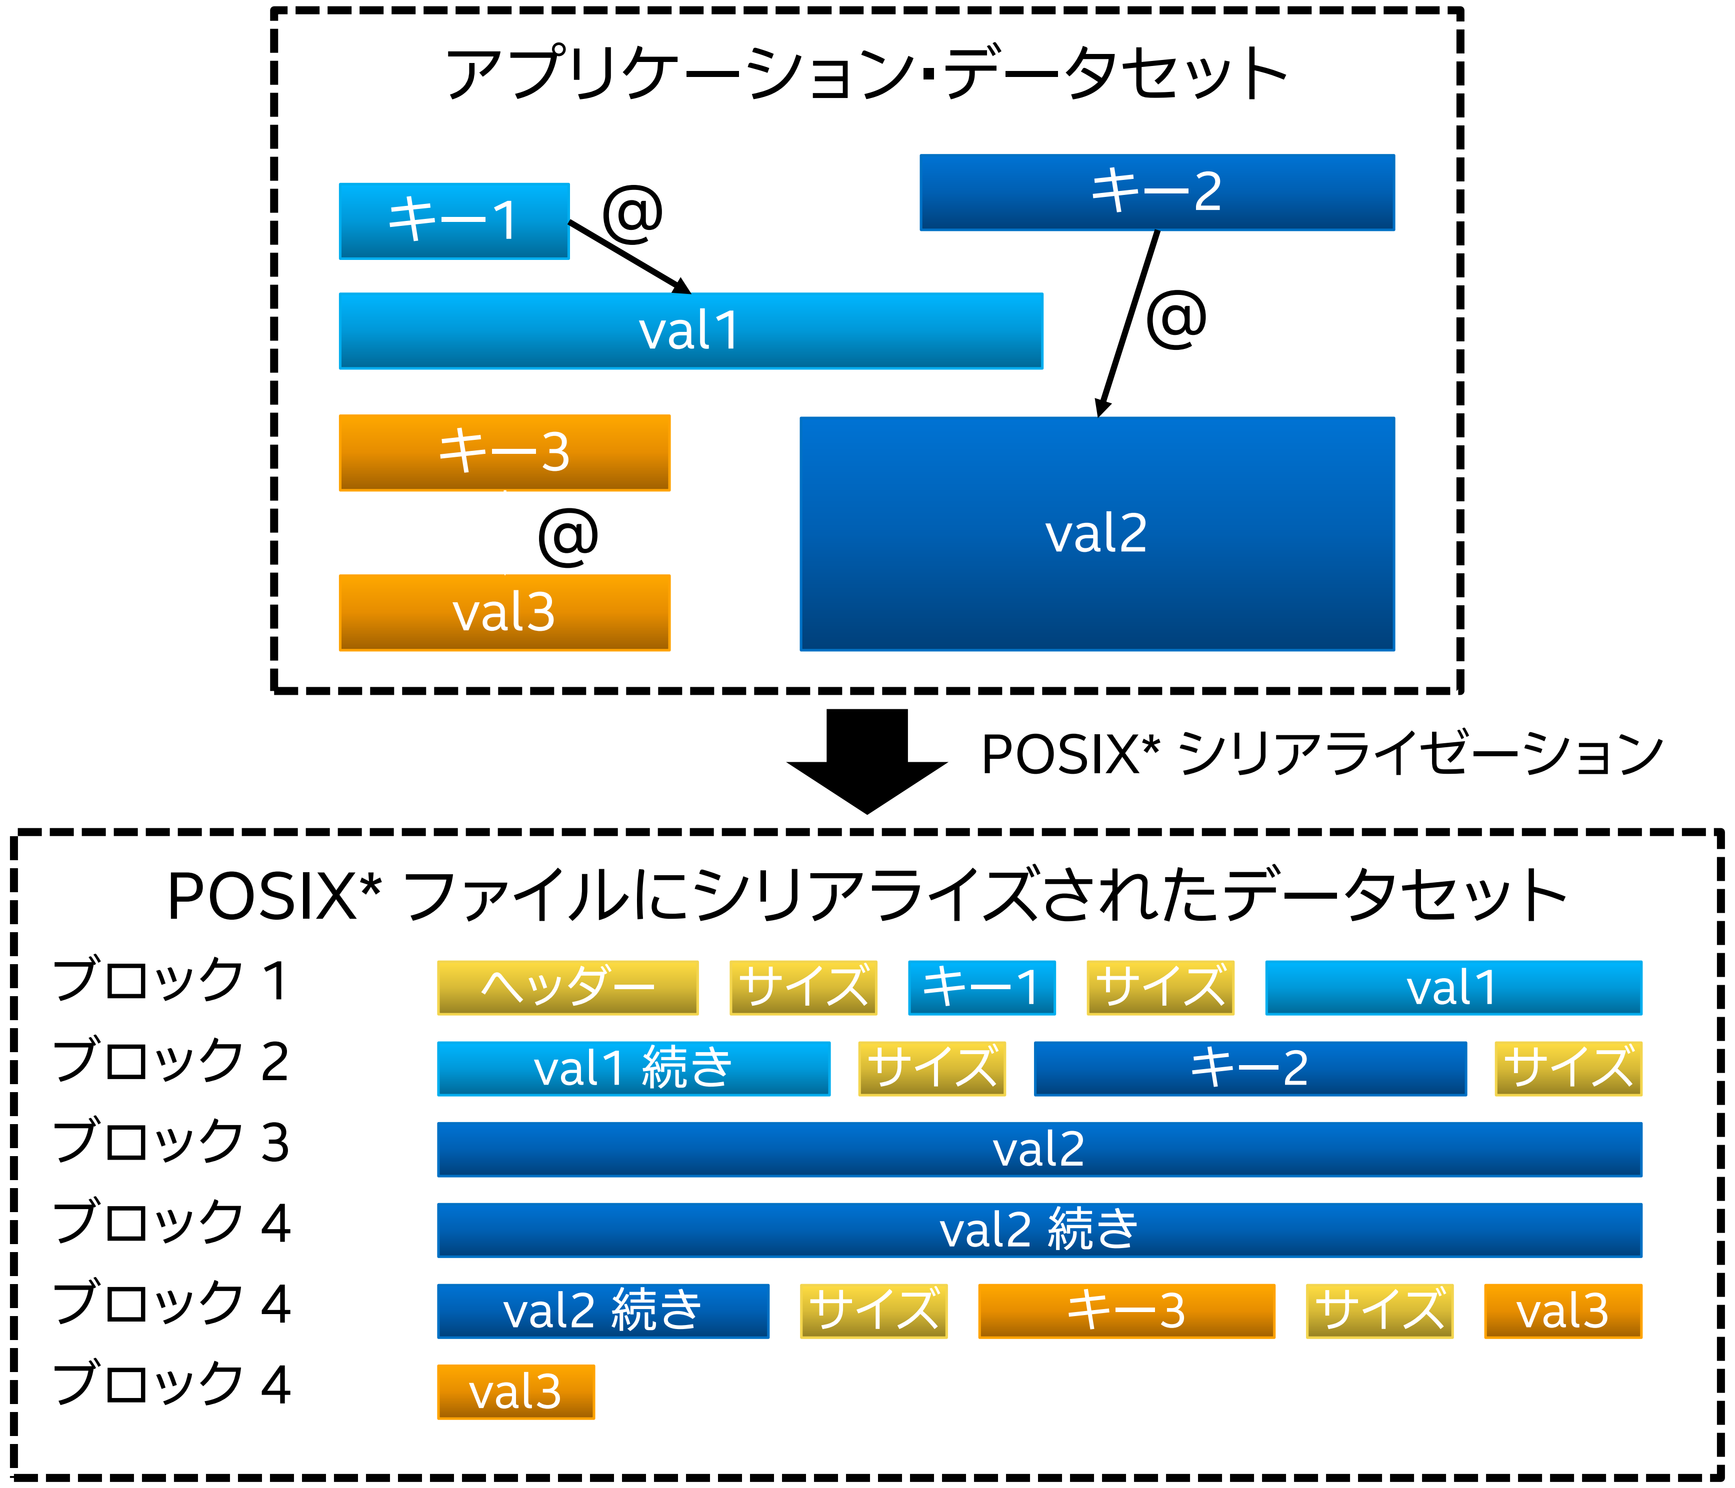

FLOPS:秒あたりバイト = 100,000 : 1

Distributed Asynchronous Object Storage

(DAOS：分散型非同期オブジェクト·ストレージ)

A new open-source, high-performance storage software solution architected for DCPMM

- Small I/Os are stored in Intel Optane DC persistent memory
- Bulk I/Os go straight to the NVMe SSDs
- Built entirely in userspace

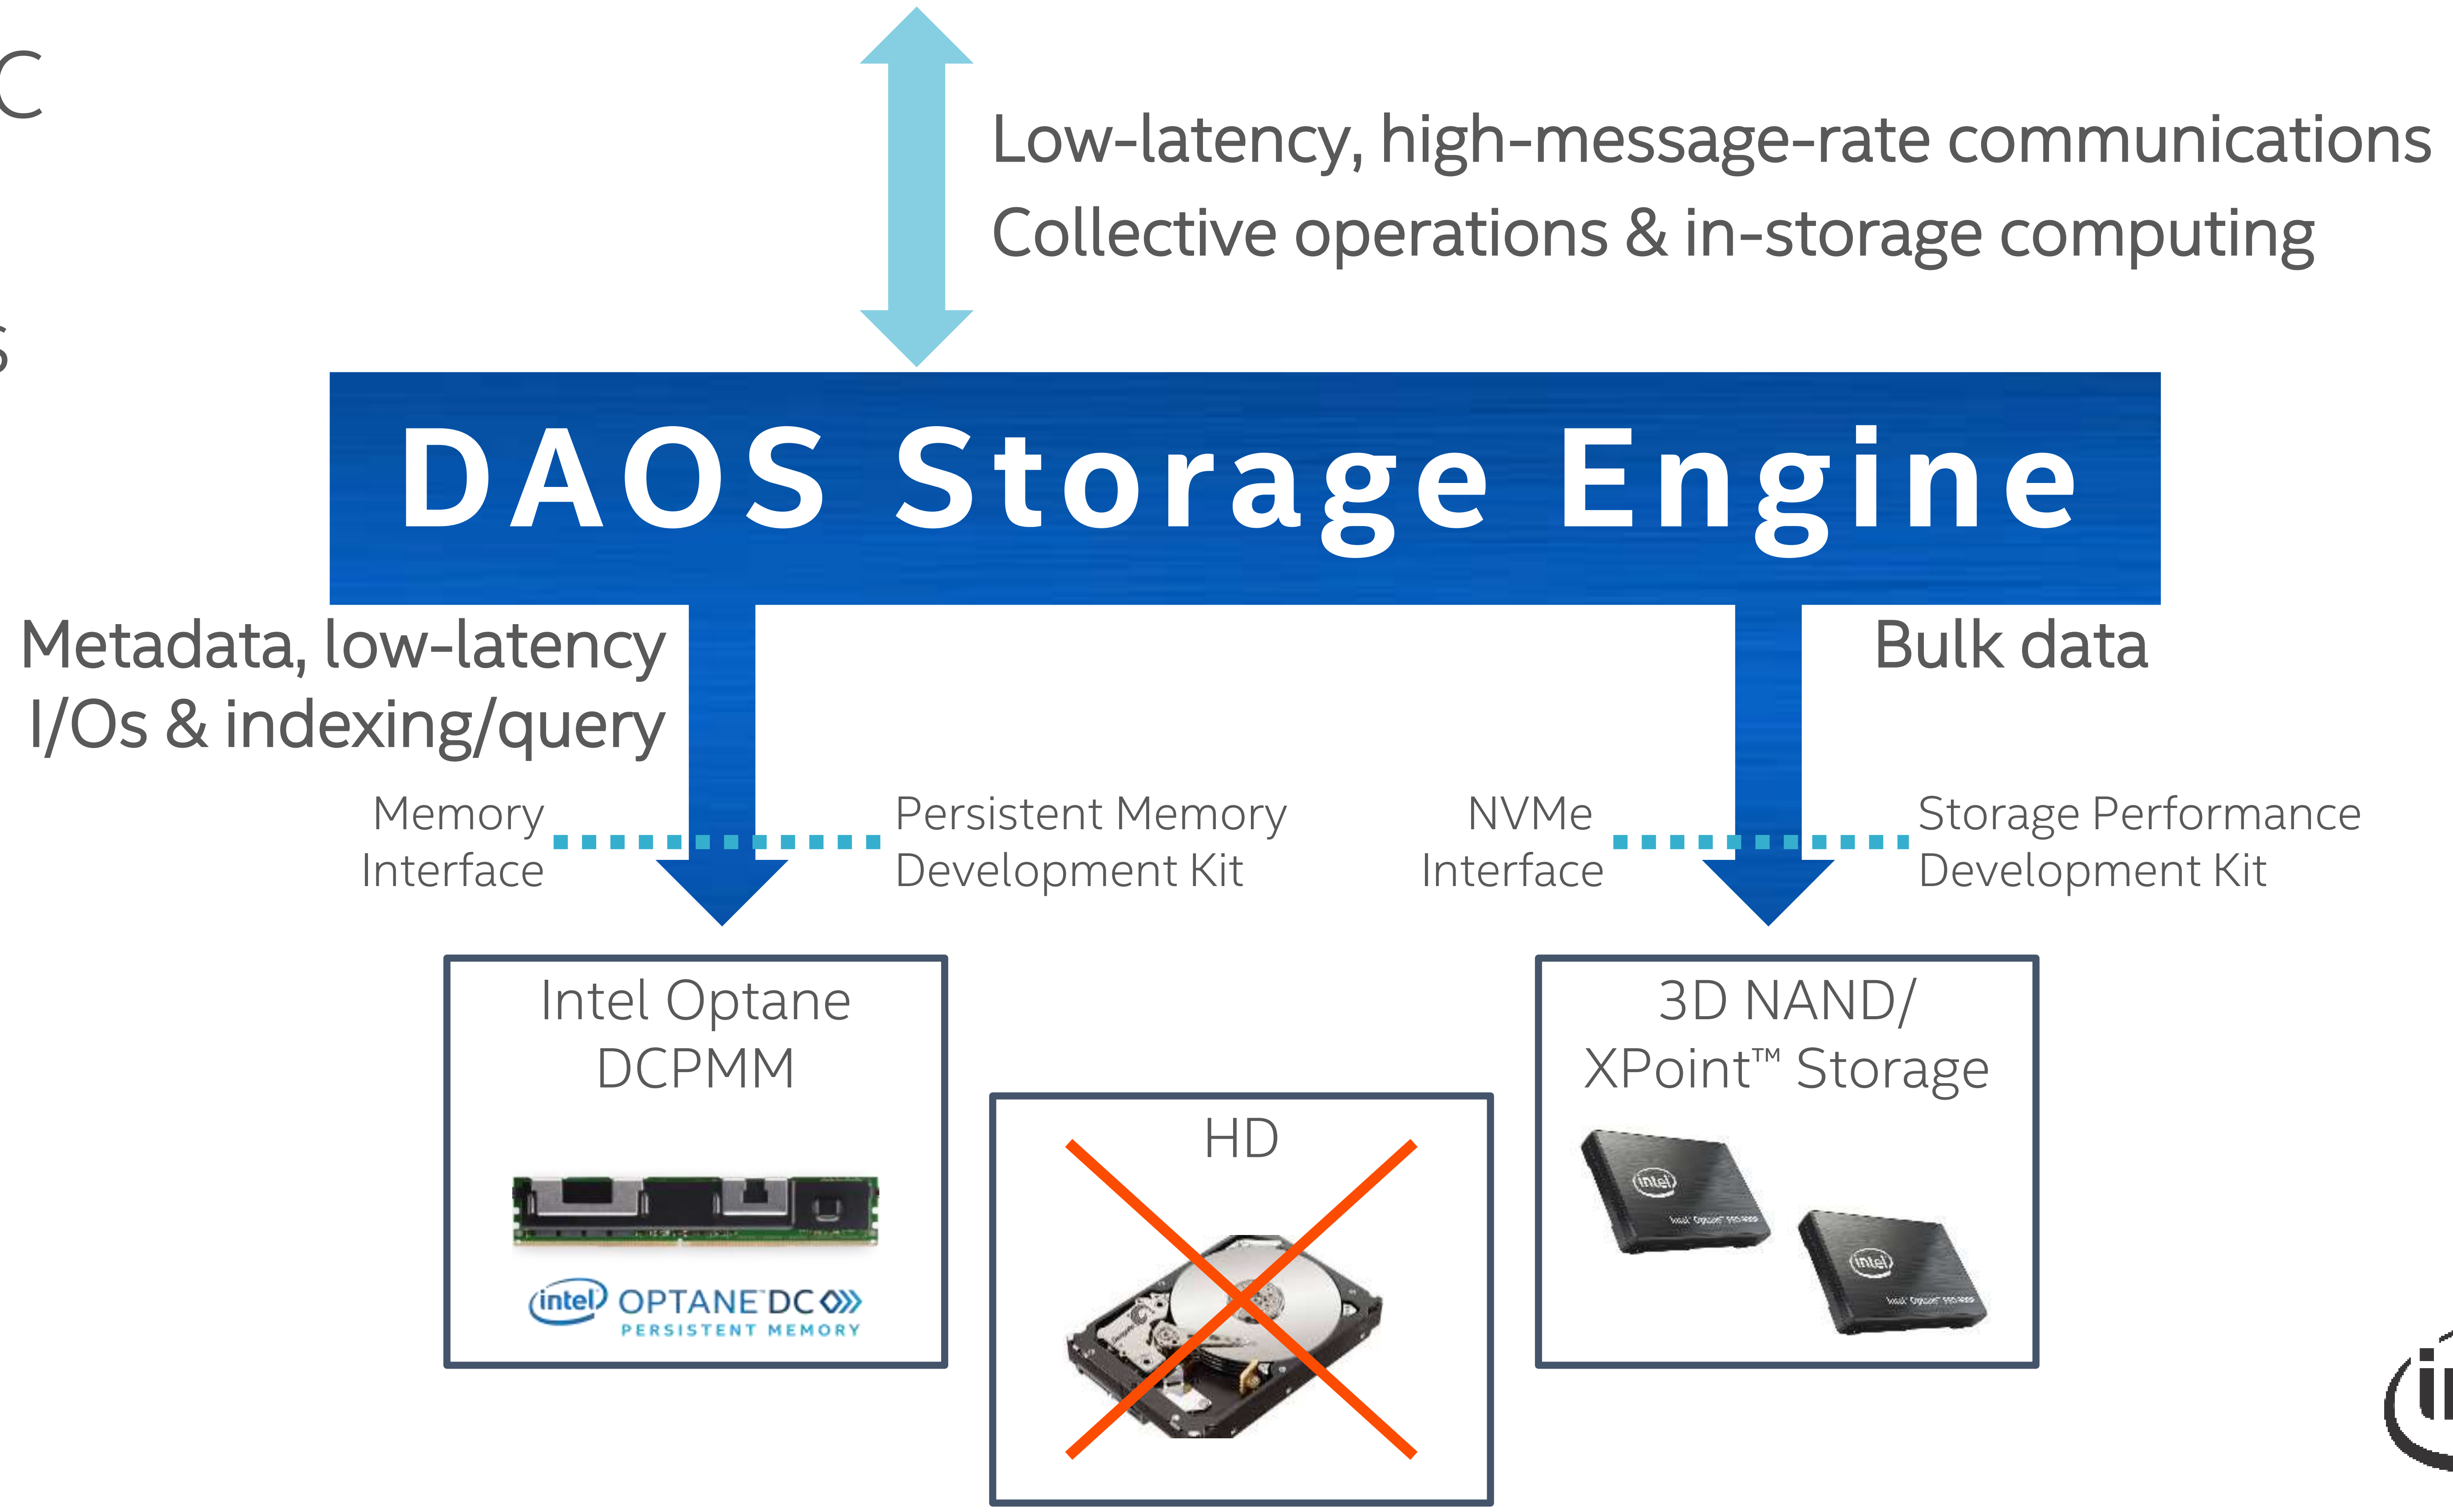

Distributed Asynchronous Object Storage

(DAOS：分散型非同期オブジェクト・ストレージ)

コンバージド・モデル

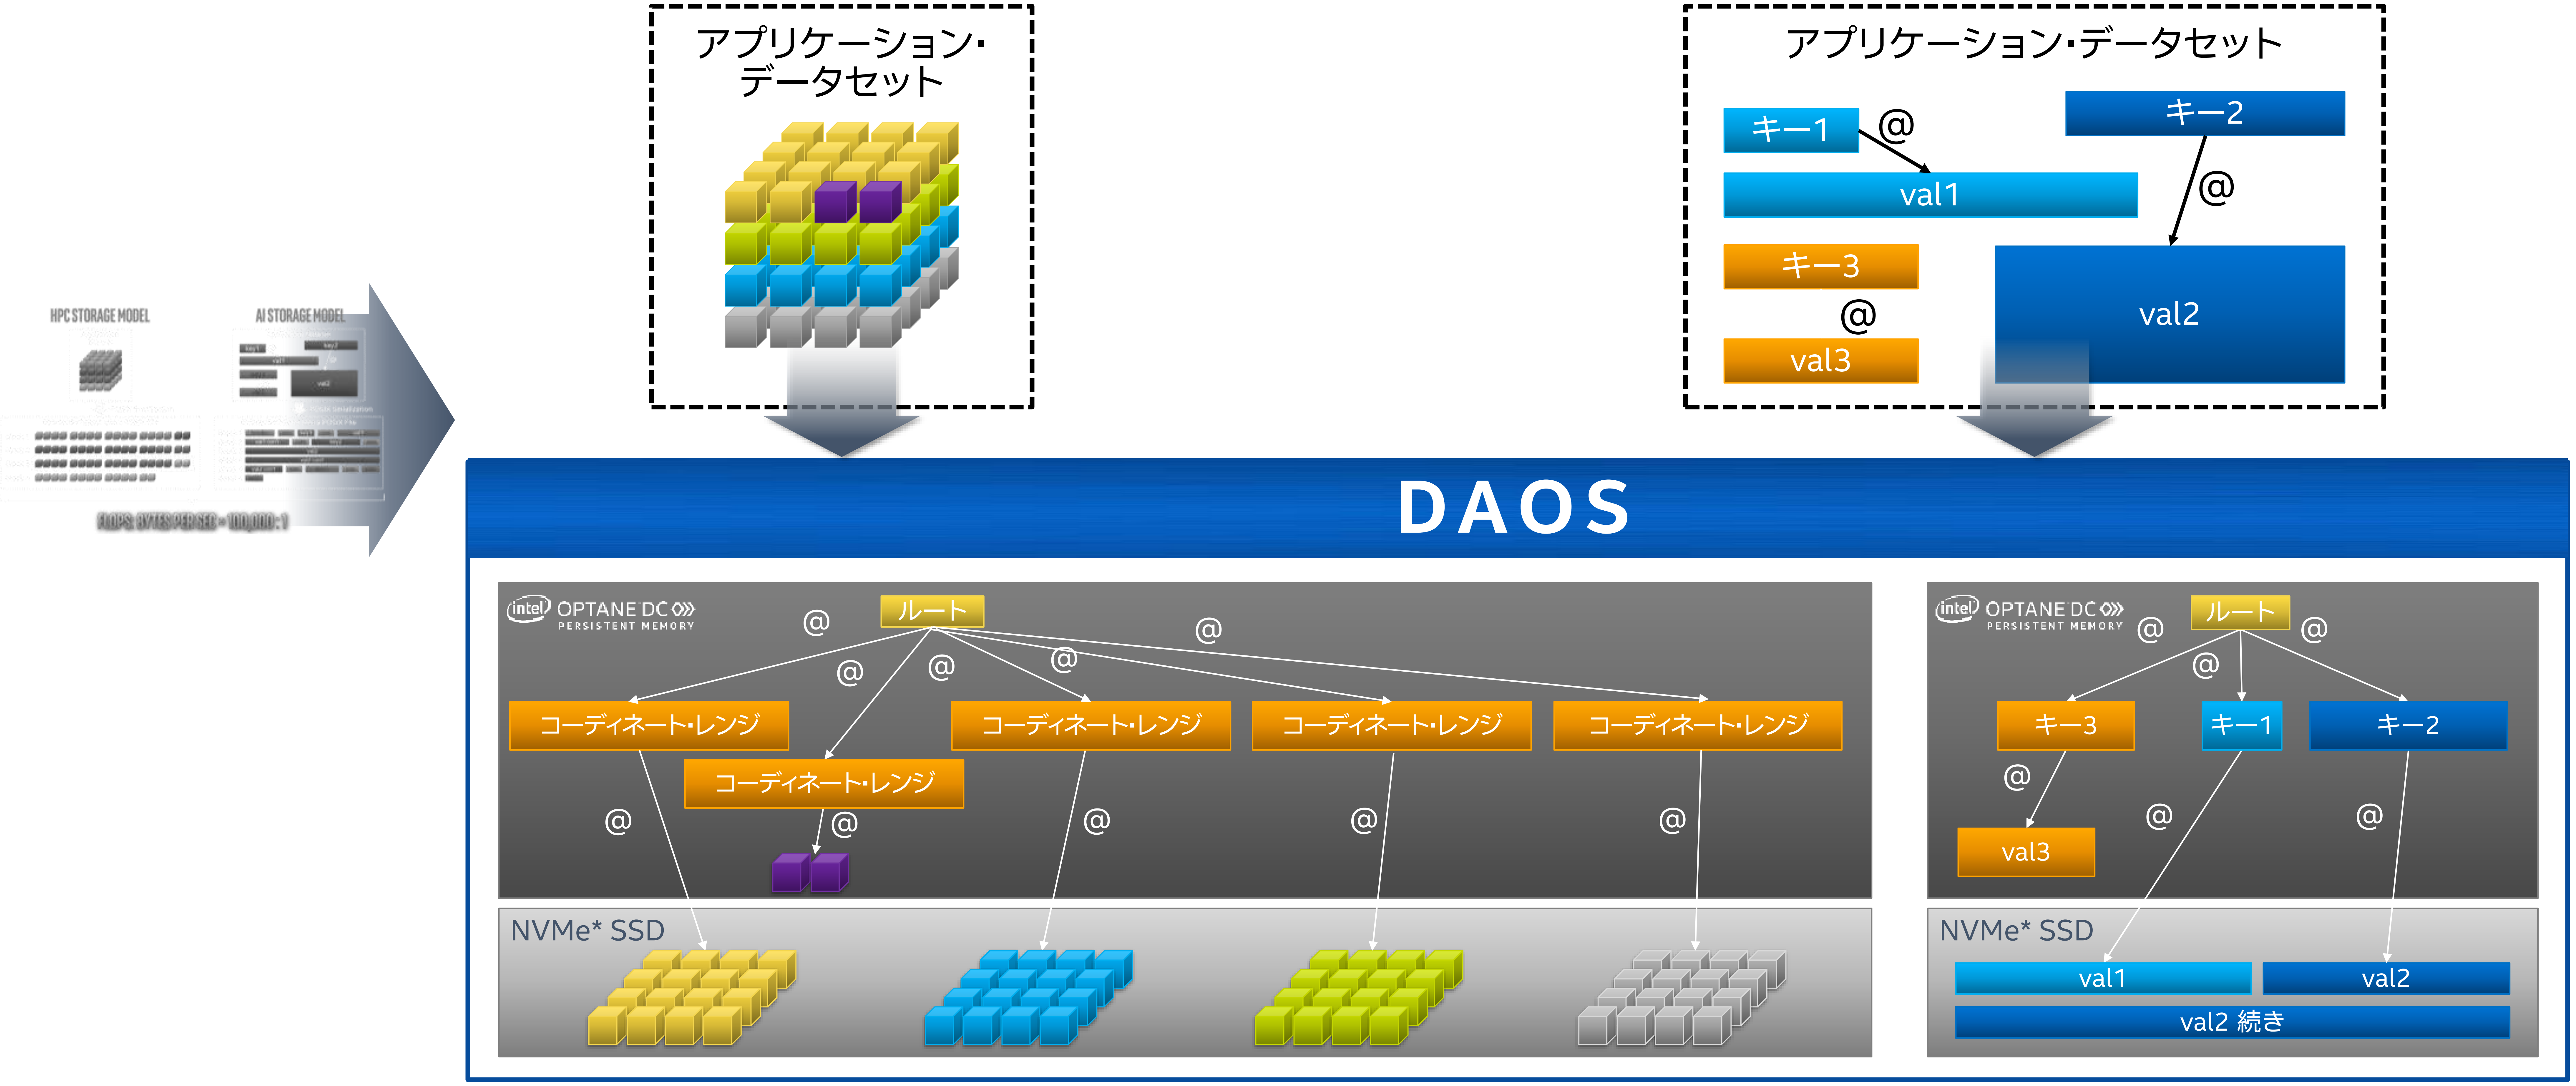

FLOPS：秒あたりバイト = 10 : 1

DAOS

GitHub にて利用可能：@ https://github.com/daos-stack/daos (英語) *

ソリューション概要は www.intel.com/hpc からダウンロード

DAOS パブリック・ロードマップ：https://wiki.hpdd.intel.com/display/DC/Roadmap (英語)

「Argonne Leadership Computing Facility は、Aurora の一部である DAOS ストレージシステムが初めて本格的に本番環境展開するもので、2021年にアメリカで始まる初のエクサスケール・システムです。DAOS のストレージシステムは、エクサスケール・レベルのマシンで I/O の拡張的ワークロードに必要なメタデータの容量使用率のレベルを提供できるように設計されています。」

Susan Coghlan

ALCF-X Project Director/Exascale Computing Systems Deputy Director

* Apache 2.0 ライセンスに基づき利用可能

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. Check with your system manufacturer or retailer or learn more at intel.com.

Performance results are based on testing as of the dates shown in configurations 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 information, go to 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 and SSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel.

Microprocessor-dependent optimizations in this product are intended for use with Intel® microprocessors. Certain optimizations not specific to Intel® microarchitecture are reserved for Intel® microprocessors. Please refer to the applicable product user and reference guides for more information regarding the specific instruction sets covered by this notice. Notice revision #20110804.

Cost reduction scenarios described are intended as examples of how a given Intel®-based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.
Intel does not control or audit third-party benchmark data or the websites referenced in this document. You should visit the referenced website and confirm whether referenced data are accurate.

THANK YOU!

RAJA KODURI

Business Forecast: Statements in this document that refer to Intel's plans and expectations for the quarter, the year, and the future, are forward-looking statements that involve a number of risks and uncertainties. A detailed discussion of the factors that could affect Intel's results and plans is included in Intel's SEC filings, including the annual report on Form 10-K.

APPENDIX A - INCREASING DENSITY: SSD FORM FACTOR INNOVATION

Footnote 1. As measured by Office 365* application launch with background activity (18 GB file copy). Configuration: CPU: Intel® Core™ i5-8265U CPU (4 cores, 8 threads, 1.6 GHz base frequency, 3.9 GHz max turbo frequency) on HP Envy x360 2-in-1 15.6" 15M-DR0011DX (BIOS F.03) with Intel® UHD 620 graphics and Intel® Optane™ Memory H10 (512 GB) vs. AMD* Ryzen 7 3700U CPU (4 cores, 8 threads, 2.3 GHz base frequency, 4.0 GHz max turbo frequency) on HP Envy x360 2-in-1 15.6" 15M-DS0012DX (BIOS F.07) with Radeon* Vega 10 graphics and Toshiba XG5 (512 GB), both with 8 GB DDR4 RAM. Storage Driver: Intel® Rapid Storage Technology 17.2.0.1009 for H10, Windows inbox driver for XG5. OS: Windows* 10 RS5 Version 1809, Build 17763. MS Office 365 Version 1902 Build 11328.

Footnote 2. As measured by Path of Exile* game launch with background activity (18GB local file copy), comparing AMD Ryzen* 7 3700U on HP Envy x360 2-in-1 15.6" 15M-DS0012DX (BIOS F.07) with Toshiba XG5 (TLC NAND SSD) 512GB vs. Intel® Core i7-8565U on HP Envy x360 2-in-1 15.6" 15M-DR0012DX (BIOS F.03) with Toshiba XG5 512GB vs. Intel® Core™ i7-8565U with Intel® Optane™ Memory H10 512GB; all configs with 8GB RAM. H10 configs tested with RST driver 17.2.0.1009; XG5 configs with Windows inbox driver. All configs used Windows 10 Home 64-bit version 1809, build 17763. Path of Exile version 3.6.6c.

Footnote 3. As measured by Photoshop CC 2019 and GIMP application launch with background activity (18 GB file copy). Configuration: CPU: Intel® Core™ i5-8265U CPU (4 cores, 8 threads, 1.6 GHz base frequency, 3.9 GHz max turbo frequency) on HP Envy x360 2-in-1 15.6" 15M-DR0011DX (BIOS F.03) with Intel® UHD 620 graphics and Intel® Optane™ Memory H10 (512 GB) vs. AMD* Ryzen 7 3700U CPU (4 cores, 8 threads, 2.3 GHz base frequency, 4.0 GHz max turbo frequency) on HP Envy x360 2-in-1 15.6" 15M-DS0012DX (BIOS F.07) with Radeon* Vega 10 graphics and Toshiba XG5 (512 GB), both with 8 GB DDR4 RAM. Storage Driver: Intel® Rapid Storage Technology 17.2.0.1009 for H10, Windows inbox driver for XG5. OS: Windows* 10 RS5 Version 1809, Build 17763. Photoshop CC 2019 Version 20.0.4. GIMP Version 2.10.8.

Footnote 4. As measured by PCMark 10 Standard Benchmark (App Start-up score). Configuration: CPU: Intel® Core™ i5-8265U CPU (4 cores, 8 threads, 1.6 GHz base frequency, 3.9 GHz max turbo frequency) on HP Envy x360 2-in-1 15.6" 15M-DR0011DX (BIOS F.03) with Intel® UHD 620 graphics and Intel® Optane™ Memory H10 (512 GB) vs AMD* Ryzen 7 3700U CPU (4 cores, 8 threads, 2.3 GHz base frequency, 4.0 GHz max turbo frequency) on HP Envy x360 2-in-1 15.6" 15M-DS0012DX (BIOS F.07) with Radeon* Vega 10 graphics and Toshiba XG5 (512 GB), both with 8 GB DDR4 RAM. Storage Driver: Intel® Rapid Storage Technology 17.2.0.1009 for H10, Windows inbox driver for XG5. OS: Windows* 10 RS5 Version 1809, Build 17763. PCMark 10 GUI version 2.0.2115. SystemInfo version 5.18.705. System benchmarks version 1.1.

Footnote 5. Based on Iometer testing at QD1 using a 4k random read scenario. Configuration: CPU: Intel® Core™ i5-8265U (on HP Envy x360 2-in-1 15.6" 15M-DR0011DX, BIOS F.03) vs. AMD Ryzen* 5 3500U (on HP Envy x360 2-in-1 15.6" 15M-DS0011DX, BIOS F.07); storage Intel® Optane™ Memory H10 512GB vs. Toshiba XG5 512GB (TLC NAND); both with 8GB RAM. Storage Driver: Intel® Rapid Storage Technology 17.2.0.1009 for H10, Windows inbox driver for XG5. OS: Windows* 10 Home 64-bit Version 1809, Build 17763.

APPENDIX B

Intel is Leading the Way with NVM Technology

1. 1st to 3Xnm (34nm)- https://phys.org/news/2009-07-intel-industry-nanometer-nand-solid-state.html
2. 1st to 2Xnm (25nm)- https://www.intel.com/pressroom/archive/releases/2010/20100201comp.htm
3. 1st 128GB with 1st integrated Hi-K Metal Gate Stack - https://www.pcmag.com/article2/0,2817,2397287,00.asp
4. Highest Density 3D NAND – based on launch on March 26, 2015 comparing to other NAND die in production at that time https://newsroom.intel.com/news-releases/micron-and-intel-unveil-new-3d-nand-flash-memory/
5. Areal Density. Source – IEEE. Comparing areal density of Intel measured data on 512 Gb Intel 3D NAND to representative competitors based on 2017 IEEE International Solid-State Circuits Conference papers citing Samsung Electronics and Western Digital/Toshiba die sizes for 64-stacked 3D NAND component.
6. 1st to 64 layer TLC - http://www.storagereview.com/intel_shows_off_new_tech_ships_1st_64layer_3d_nand_for_data_center
7. Source: Intel. 1st PCIe* Intel QLC 3D NAND SSD. Based on Intel achieving PRQ status of Intel® SSD D5-P4320 on 13 July 2018.
8. 1st PCIe* QLC SSD for Client - https://www.tomshardware.com/reviews/intel-ssd-660p-qlc-nvme,5719.html