intel® XEON®

Better Performance for Single Cell Genomics with Intel® Xeon® Scalable Processors[1]

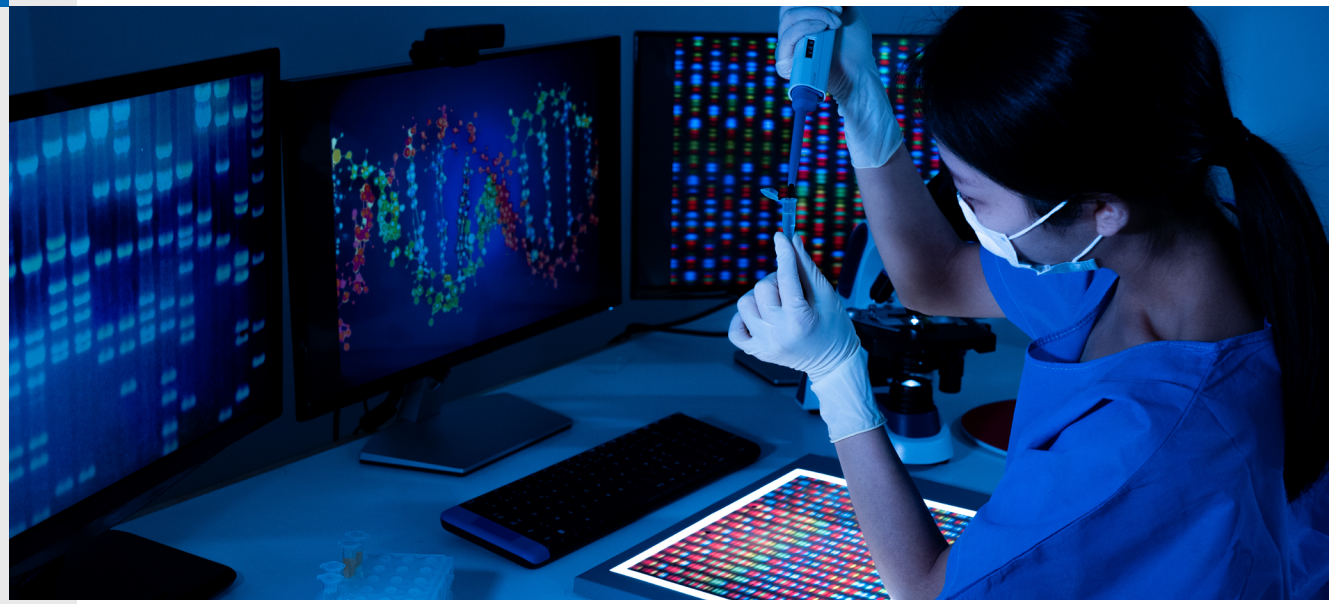

We live in an era of digital biology where artificial intelligence (AI) and high-performance computing (HPC), combined with biological sciences, are transforming the world. For example, researchers in genomics processing rely on AI and HPC methods to design more effective drugs with fewer side effects. Single-cell sequencing and the subsequent data analysis help researchers understand differences in the behavior of cells in the same organ/tissue to offer insight into how they respond to certain diseases and drugs. Cells found in an organ have the same DNA but may not have the same RNA content. Single-cell sequencing looks at this RNA content for the cells in an organ to explain why some seemingly identical cells behave differently. Such real-life workflows require acceleration of end-to-end data processing and AI algorithms (e.g., dimensionality reduction, clustering, visualization) for analysis of the single cell sequencing data. For these types of analyses, CPUs provide leading performance for HPC and AI end-to-end workflows for both on-premises and the cloud.[1]

1.3M Single Cell Genomics Analysis

Time to Complete Analysis Measured in Seconds (***Lower is Better***)

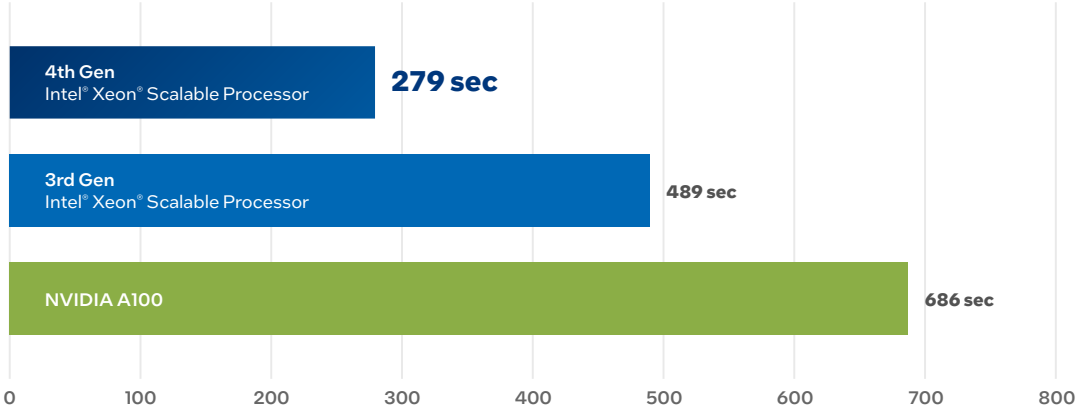

Faster End-to-End Machine Learning Performance[1,2]

3rd Gen Intel® Xeon® Scalable processors on software optimized for CPUs perform up to **1.4x** faster than NVIDIA A100 GPUs[1]

4th Gen Intel® Xeon® Scalable processors on software optimized for CPUs perform up to **2.5x** faster than NVIDIA A100 GPUs.[2]

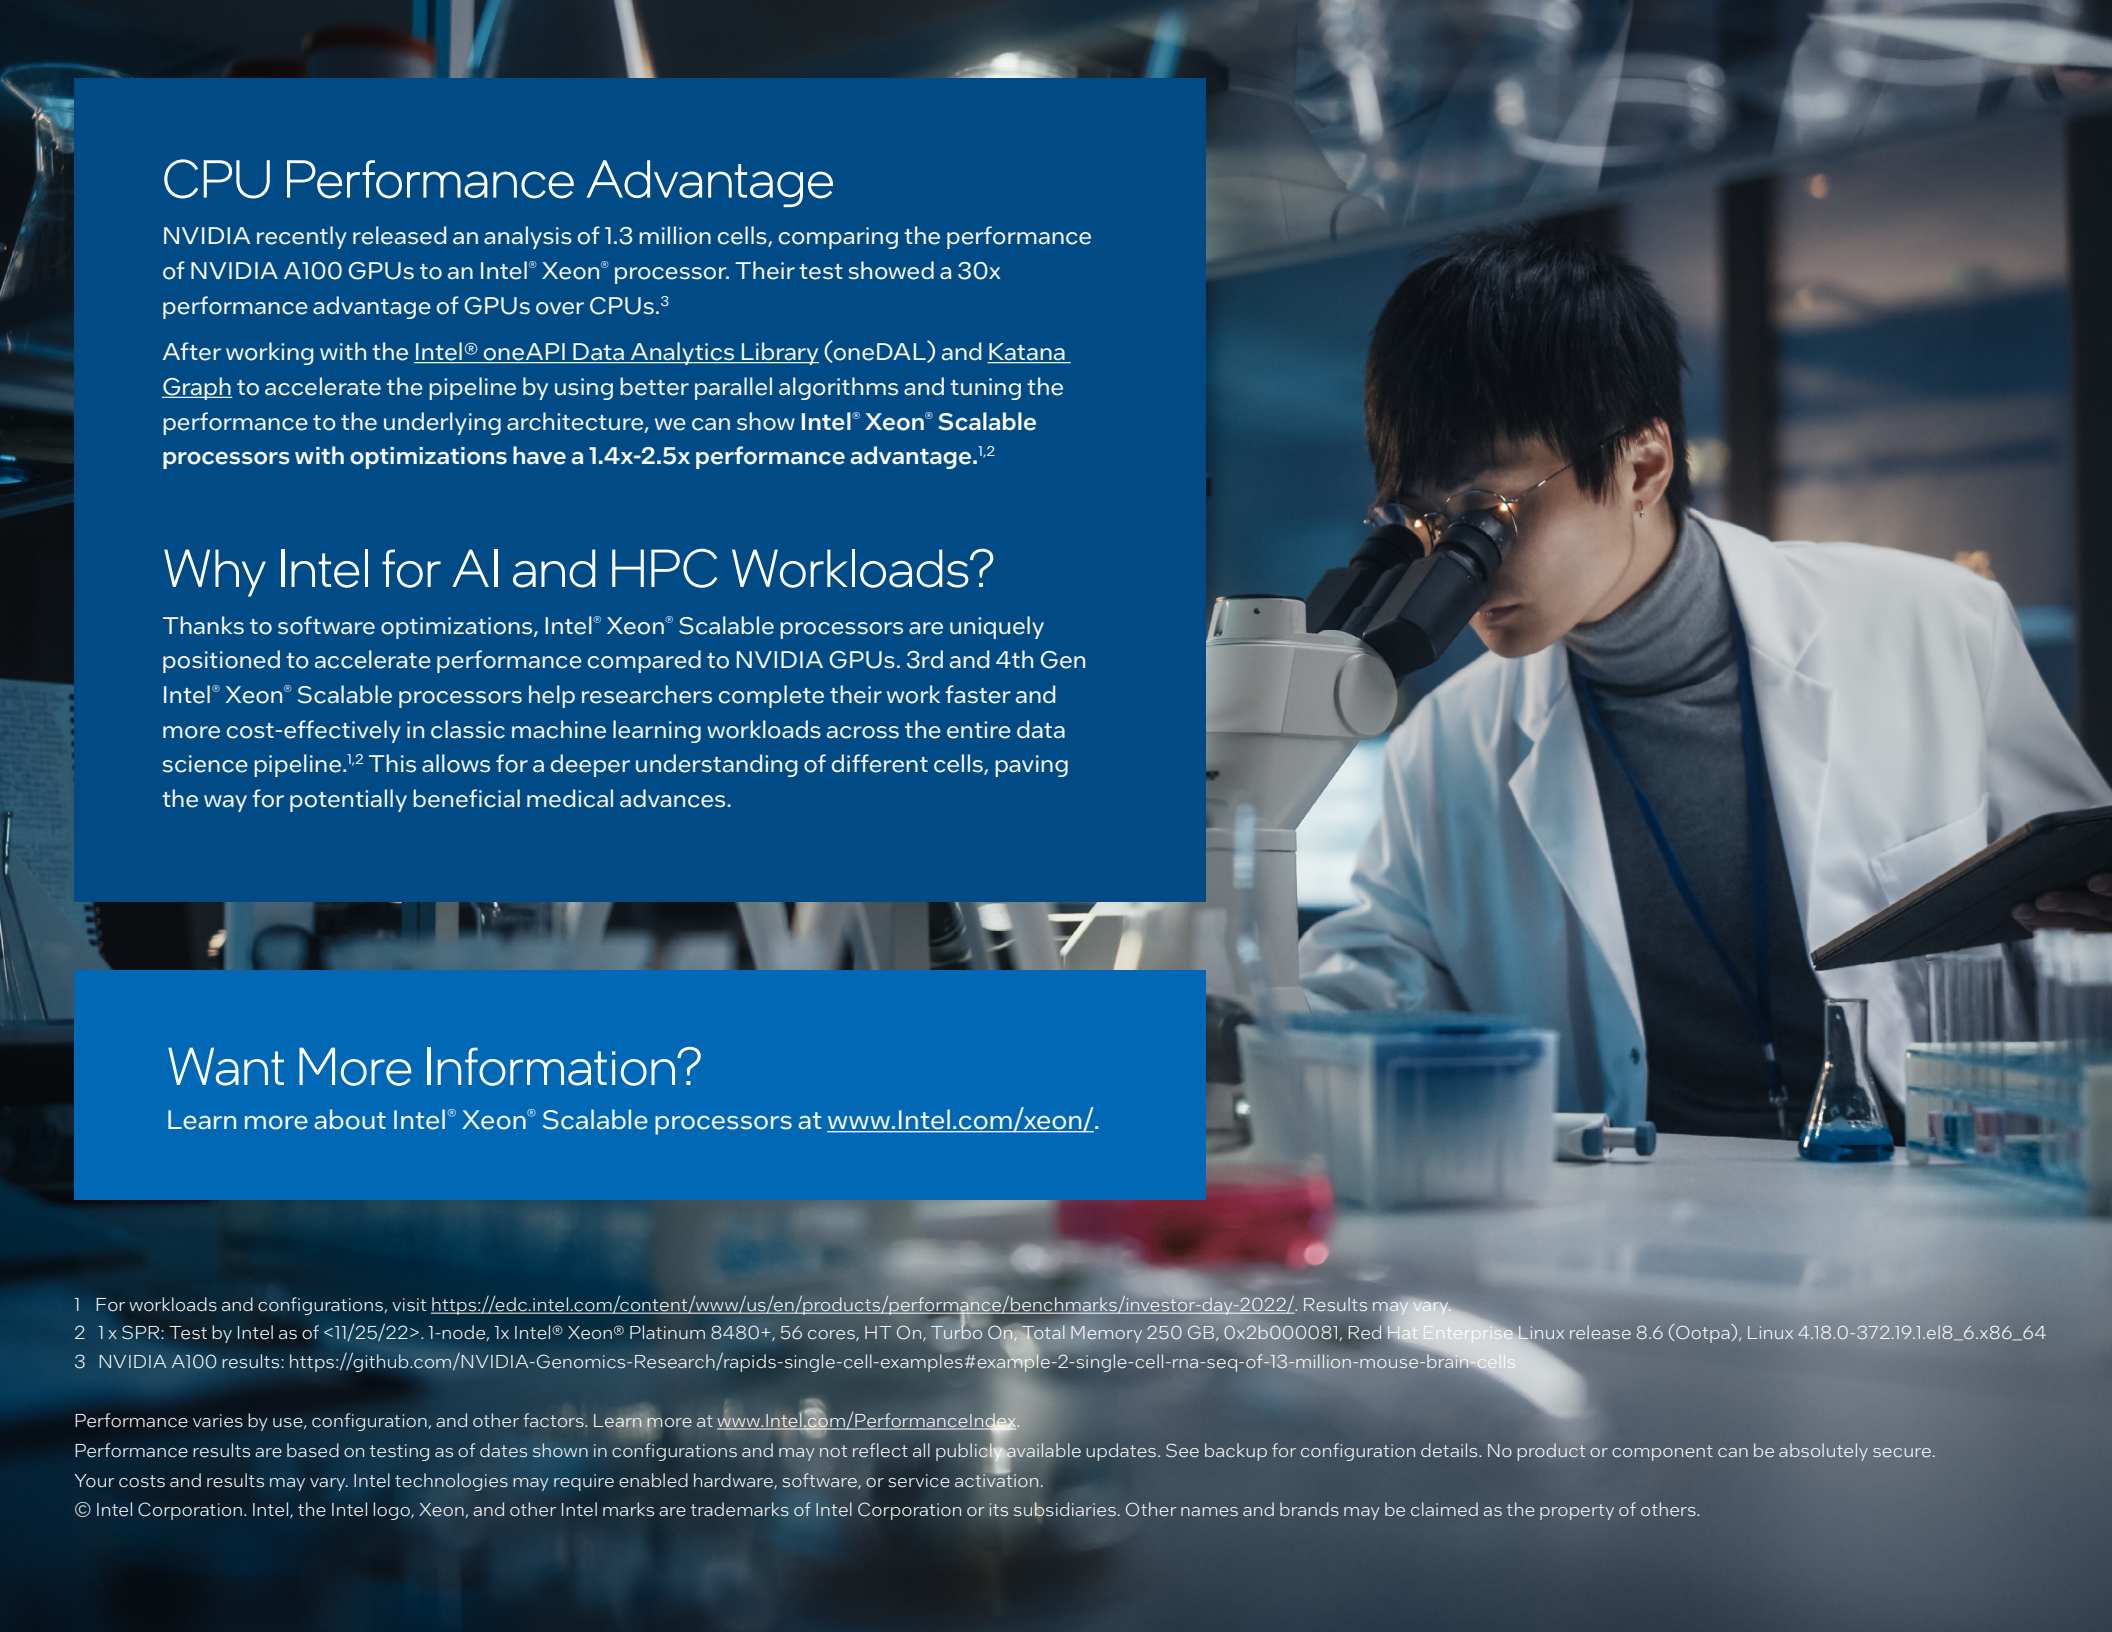

CPU Performance Advantage

NVIDIA recently released an analysis of 1.3 million cells, comparing the performance of NVIDIA A100 GPUs to an Intel® Xeon® processor. Their test showed a 30x performance advantage of GPUs over CPUs.[3]

After working with the Intel® oneAPI Data Analytics Library (oneDAL) and Katana Graph to accelerate the pipeline by using better parallel algorithms and tuning the performance to the underlying architecture, we can show **Intel® Xeon® Scalable processors with optimizations have a 1.4x-2.5x performance advantage.**[1,2]

Why Intel for AI and HPC Workloads?

Thanks to software optimizations, Intel® Xeon® Scalable processors are uniquely positioned to accelerate performance compared to NVIDIA GPUs. 3rd and 4th Gen Intel® Xeon® Scalable processors help researchers complete their work faster and more cost-effectively in classic machine learning workloads across the entire data science pipeline.[1,2] This allows for a deeper understanding of different cells, paving the way for potentially beneficial medical advances.

Want More Information?

Learn more about Intel® Xeon® Scalable processors at www.Intel.com/xeon/.

1 For workloads and configurations, visit https://edc.intel.com/content/www/us/en/products/performance/benchmarks/investor-day-2022/. Results may vary.
2 1x SPR: Test by Intel as of <11/25/22>. 1-node, 1x Intel® Xeon® Platinum 8480+, 56 cores, HT On, Turbo On, Total Memory 250 GB, 0x2b000081, Red Hat Enterprise Linux release 8.6 (Ootpa), Linux 4.18.0-372.19.1.el8_6.x86_64
3 NVIDIA A100 results: https://github.com/NVIDIA-Genomics-Research/rapids-single-cell-examples#example-2-single-cell-rna-seq-of-13-million-mouse-brain-cells

Performance varies by use, configuration, and other factors. Learn more at www.Intel.com/PerformanceIndex.
Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.
Your costs and results may vary. Intel technologies may require enabled hardware, software, or service activation.