Technical Research Report

Benefits and Comparison of Intel® Xeon® Processors with E-Cores, P-Cores, and AP-Cores for Select Workloads

This report evaluates the benefits of workload targeting and distribution across different Intel Xeon processor types, namely Efficient-cores (E-cores), Performance-cores (P-cores), and Advanced Performance–cores (AP-cores).

Executive Summary

As modern data centers expand to support AI and other data-driven services, it's increasingly critical to select the right processor architecture. Intel Xeon processors offer three specialized core types: Efficient-cores (E-cores), Performance-cores (P-cores), and Advanced Performance–cores (AP-cores). Each core type is optimized for specific workload patterns, operational priorities, and IT use cases. Each core type also offers different specifications and performance.

This paper shows that matching workloads to the most suitable core types enhances performance and lowers energy consumption for businesses. Therefore, by adopting hybrid Intel Xeon processor–based architectures, enterprises can balance performance with cost and energy efficiency across diverse application needs.

In this study, Prowess Consulting evaluated how each core type behaves when running various applications that represent the most common workloads associated with each core type:

- WRK® web benchmark for web concurrency (E-cores)
- ResNet®-50 for AI inference (P-cores)
- Nanoscale Molecular Dynamics (NAMD) for high-performance simulation (AP-cores)[1]

The goal of our testing was to determine how performance and energy consumption vary when workloads are purposefully distributed or targeted to the various core types.

Our conclusions show that:

- E-cores deliver strong performance per watt and high density for highly parallel services
- P-cores provide responsive and predictable performance for latency-sensitive analytics and transactional workloads
- AP-cores deliver sustained throughput for large, multithreaded scientific and AI workloads

Highlights

Prowess Consulting found the following workload impacts across E-cores, P-cores, and AP-cores:

WEB-SERVING THROUGHPUT (WRK2): Systems with E-cores achieved throughput within **0.25%** of those with P-cores.

RESNET-50: Systems with AP-cores delivered up to **120%** higher CPU-based AI inference throughput than those with P-cores.

NAMD: Systems with AP-cores completed **84%–87%** more molecular simulation work than those with P-cores by running more jobs concurrently.

Evaluating Core Performance

As workload volumes and variety grow, shifting how and where they run can have an incremental impact on performance and costs. For that reason, it's important to consider what kinds of workloads your business will be running and the best processors on which to run those workloads. That's why this study examines how Intel Xeon processors with E-cores, P-cores, and AP-cores respond under three distinct workload categories that represent common enterprise use cases:

- **WRK2 web benchmark** to evaluate concurrency and efficiency
- **ResNet-50** to measure inference and AI responsiveness
- **NAMD** to measure large-scale simulation and sustained multithreading

Within each of these applications, we measured five metrics: performance, energy efficiency, thermal behavior, scheduling considerations, and system-level impacts. Together, these metrics provide detailed insights into how the application is performing and how the processor handles the workload.

Why Intel Xeon Processors

Modern IT environments support a mix of latency-sensitive applications, background services, analytics pipelines, AI inference tasks, and compute-heavy simulations. Because each of these workloads places varying demands on system hardware, hybrid Intel Xeon processor–based architectures present a good option for matching each workload to the most suitable core type.

Hybrid Intel Xeon processor–based architectures combine E-cores, P-cores, and AP-cores to balance workloads with the most appropriate compute resources. This allows users to dynamically balance workloads by placing latency-sensitive or bursty tasks on P-cores, parallel or background tasks on E-cores, and compute-intensive operations on AP-cores. When organizations assign workloads to the most suitable processor cores, they can enhance system responsiveness and efficiency while minimizing the need for manual adjustments.

These architectures also make it easier to run diverse workloads on a single server platform. With heterogeneous cores and built-in accelerators, the same infrastructure can support AI inference, analytics, general compute, virtualization, high-performance computing (HPC) workloads, and edge services.

Business Value

For a business, mapping workloads to the right type of core can help improve infrastructure performance and avoid placing superfluous demands on ill-matched hardware, which could force unnecessary scale-out deployments. When the right task runs on the right core, you can experience fewer bottlenecks, more predictable latency, and better overall throughput (as measured by requests per second [RPS]). At scale, this optimization can lead to improved service levels and more efficient use of resources.

Understanding E-Cores, P-Cores, and AP-Cores

Each of these core types excels in different operational scenarios, and businesses increasingly deploy them in targeted ways.

E-Cores: High-Density and Background Tasks

E-cores are designed for parallel work and to provide higher performance per watt. Processors with E-cores are ideal for workloads with high concurrency but modest per-thread demands. These cores offer lower single-thread performance but superior scalability and efficiency. They thus help organizations meet sustainability goals and comply with emerging regulatory requirements related to power consumption and carbon emissions. Specifically, they can improve virtual machine (VM)/container density, reducing the physical footprint and cooling requirements of data center deployments.

Common business applications include:

- Front-end web services, microservices, and lightweight API endpoints
- Background or maintenance tasks, such as logging, telemetry processing, and batch jobs
- Virtualized environments with large VM counts
- Kubernetes® clusters running high-density container workloads
- Edge deployments where power, cooling, and physical space are limited

P-Cores: Latency-Sensitive and Bursty Workloads

P-cores are designed for high-performance, low-latency workloads, and they offer strong single-thread speeds. Processors with P-cores are ideal for analytics and real-time decision-making, and they offer more performance per core, but at a higher power draw than processors with E-cores. Businesses typically assign these cores for:

- Real-time AI inference and analytics
- Transaction-heavy online transaction processing (OLTP) databases
- Retail intelligence systems and real-time decision engines
- Customer-facing applications that must respond instantly
- Augmented reality (AR)/virtual reality (VR), rendering, and other compute-intensive foreground tasks

AP-Cores: AI, HPC, and Simulation Workloads

AP-cores are designed for large-scale, memory-intensive workloads. Processors with AP-cores are ideal for heavy compute loads with high bandwidth and enhanced parallelism. These processors are often part of a larger, highly capable system with extra hardware meant to accelerate AI, compute, or data movement workloads. Common uses for AP-cores include:

- Enterprise AI workloads using GPU-like acceleration paths
- Large-scale simulations (such as those involving NAMD, finite element analysis [FEA], or computational fluid dynamics [CFD])
- Real-time industrial analytics and embedded AI
- Memory-intensive compute clusters
- Workflows that combine general-purpose compute with NPUs, GPUs, or other accelerators

Conclusions from Testing and Analysis

Our findings show that targeted workload assignments, combined with optimizing the hardware on which they run, can significantly impact performance and efficiency. For businesses, these findings are especially relevant to the use cases discussed earlier. User experience depends more on smart workload scheduling than on raw hardware capacity. Hybrid scheduling, where the operating system (OS) or orchestrator assigns tasks to the most appropriate core type, can produce smoother performance and fewer system-wide slowdowns.

- We found that E-cores deliver superior performance-per-watt for parallel, background, and concurrency-heavy workloads. Offloading non-critical tasks to E-cores improves efficiency, reduces thermal stress, and helps maintain consistent service levels at scale. This approach supports longer system lifespans and more predictable infrastructure growth.
- P-cores deliver their strongest value under specific conditions. They provide responsive, predictable performance when applications depend on low latency or strong single-thread execution, such as transactional systems or per-task analytics. When P-cores are used for background or highly parallel workloads, however, power efficiency declines without a corresponding performance benefit.
- AP-cores are designed for workloads that require sustained, multithreaded throughput. In our testing, they delivered the highest aggregate performance for HPC simulations and CPU-based AI inference by efficiently running many jobs concurrently. AP-cores provide an efficient foundation for organizations executing continuous scientific computing, large-scale analytics, or vision-style inference workloads where total work completed matters more than individual job latency.

Test Plan and Results

We tested each application on a different server in a specific configuration to best simulate real-world scenarios. We confined each configuration and the subsequent tests to that hardware to ensure consistent data. For this testing, we used the following three server configurations:

- **Dell™ PowerEdge™ R770 server with a configuration using E-cores** (2 x Intel Xeon 6740E processor, 96 cores each)
- **PowerEdge R770 server with a configuration using P-cores** (2 x Intel Xeon 6767P processor, 64 cores each)
- **PowerEdge R770AP server with a configuration using AP-cores** (2 x Intel Xeon 6978P processor, 120 cores each)

All systems used Red Hat® Enterprise Linux® 10, identical baseline sysbench runs, and consistent networking and storage configurations (with the systems using AP-cores using higher-capacity power supplies and memory). For a complete breakdown of the configurations, refer to the **Appendix**.

Tested Workloads

By analyzing the behavior of E-cores, P-cores, and AP-cores under controlled conditions, we gained a clear understanding of how the following workloads stress a platform, how the various cores respond under sustained load, and how businesses can align their application demands with the optimal core architecture.

WRK2 Web Benchmark

WRK2 is a high-performance HTTP benchmarking tool designed to stress-test web servers using an event-driven architecture and large numbers of concurrent connections. It generates realistic traffic patterns and captures detailed metrics such as RPS, latency distributions, and socket-level errors. Because it exposes how a system behaves under sustained parallel load, WRK is well-suited for comparing heterogeneous Intel Xeon processor architectures, including those with E-cores, P-cores, and AP-cores, in unified test conditions.

Our testing focused on each core type's ability to sustain throughput, manage queue depth, and maintain stable latency under varying concurrency levels. By examining how different Intel Xeon processor architectures respond to identical WRK2 workloads, we observed clear differences in scheduling efficiency, thermal behavior, and power draw. These insights can help IT architects determine the most efficient core assignments for web services, API endpoints, and distributed application front ends, enabling data centers to optimize both performance and performance per watt at scale.

Results

The WRK2 benchmark highlights that web-serving workloads are driven more by the ability to manage large numbers of concurrent, lightweight requests than by peak single-thread performance. This behavior was consistent across all tested configurations. At low to moderate connection counts, the systems with E-cores, P-cores, and AP-cores delivered nearly identical request throughput, with meaningful differences appearing only as concurrency increased.

With up to 10,000 requests with 50 connections, the systems with E-cores and P-cores maintained throughput within 0.25% of each other. We observed a similar level of parity between the systems with E-cores and those with AP-cores at 10,000 requests with 500 connections.

The systems with E-cores performed on par with both those with P-cores and those with AP-cores across most WRK2 throughput levels. The workload itself did not demand high-performance cores to maintain acceptable service levels. CPU utilization remained high throughout testing, reflecting efficient use of available compute resources across all architectures.

Conclusions

This testing shows that systems with E-cores provide a strong and efficient match for web-serving and API-style workloads that emphasize high concurrency and lightweight request handling. WRK performance was broadly comparable across systems with E-cores, P-cores, and AP-cores, which reinforces that higher-performance cores provide a limited throughput advantage for this type of workload.

Deploying systems with P-cores or AP-cores is unlikely to yield meaningful performance gains for Apache HTTP–style workloads, and it could instead increase power consumption without clear benefit. Platforms with E-cores deliver sufficient performance while aligning more closely with efficiency and power-conscious deployment goals. This fact makes them well-suited for front-end services, microservices, and other highly parallel web workloads.

AI-Inference Workloads Using ResNet-50 with Intel-Optimized Toolkits

ResNet-50 is a widely adopted convolutional neural network. It is commonly used to evaluate real-world image-classification inference performance. In this study, we used ResNet-50 with a combination of Intel-optimized software, including:

- Intel® OpenVINO™ toolkit
- Intel® oneAPI Base and HPC toolkits
- Intel-optimized TensorFlow™

This approach allowed us to evaluate CPU-based AI inference performance across E-cores, P-cores, and AP-cores while reflecting common enterprise deployment patterns for vision-style inference workloads.

By analyzing how E-cores, P-cores, and AP-cores process identical OpenVINO toolkit–driven ResNet-50 inference tasks, we observed differences in vector execution behavior, thread-scheduling efficiency, and power draw across various batch sizes and precisions. Our testing also captured how each architecture responds to sustained inference pressure, including frequency scaling, thermal stability, and memory bandwidth sensitivity. These measurements provide a technical foundation for understanding how heterogeneous Intel Xeon processor architectures support AI inference pipelines across edge, data center, and hybrid environments.

Results

Our testing with the OpenVINO toolkit using ResNet-50 showed clear architectural separation for CPU-based AI inference workloads. Systems with AP-cores delivered the highest inference throughput and outpaced both systems with P-cores and those with E-cores across all evaluated scenarios. The performance of systems with AP-cores ranged from modestly higher than that of systems with P-cores to more than double in certain cases (see Table 1).

Table 1 | Maximum and minimum relative performance for systems with AP-cores versus systems with P-cores by precision mode and block size

Precision Mode	Block Size	Maximum Performance for Systems with AP-Cores vs. Systems with P-Cores	Minimum Performance for Systems with AP-Cores vs. Systems with P-Cores
INT8	1	119% higher	58% higher
INT8	16	113% higher	52% higher
FP32	1	120% higher	45% higher
FP32	16	101% higher	26% higher

However, this higher raw performance tells only part of the story. Systems with P-cores ran the ResNet-50 workload reliably and produced stable, interpretable results across a wide range of configurations. They also consumed less power, both with their processors using less energy and by having less memory drawing energy. The lower power-draw of systems powered by P-cores means that the performance per watt of such systems could exceed that of systems with AP-cores (see Table 2). 44% of the combinations of precision mode and block size for ResNet-50 tested on systems powered by P-cores were at parity with or exceeding the performance per watt of those of AP-core systems.[2] This means that higher throughput with AP-cores doesn't necessarily translate into better efficiency, and the lower power draw with systems with P-cores can be important to organizations looking to reduce their total power consumption.

Table 2 | Maximum and minimum relative performance per watt for systems with P-cores versus systems with AP-cores by precision mode and block size

Precision Mode	Block Size	Maximum Performance per Watt for Systems with P-Cores vs. Systems with AP-Cores	Minimum Performance per Watt for Systems with P-Cores vs. Systems with AP-Cores
INT8	1	52% higher	56% higher
INT8	16	132% higher	53% higher
FP32	1	132% higher	52% higher
FP32	16	209% higher	42% higher

We evaluated systems with E-cores more selectively for ResNet-50 inference. While capable of executing the workload, these systems consistently delivered lower throughput than both those with AP-cores and those with P-cores. This outcome aligns with architectural tradeoffs that favor efficiency rather than per-instance inference performance.

We also examined scaling behavior across multiple configurations, including INT8 and FP32 precision modes, varying block sizes, and different core counts. In some scenarios, lower core counts resulted in higher throughput. Despite these variations, the overall performance ordering remained unchanged: systems with AP-cores led, followed by those with P-cores, with systems with E-cores trailing.

Conclusions

Systems with AP-cores provide the strongest results for ResNet-50 inference on CPUs using the OpenVINO toolkit. They consistently achieved the highest throughput across configurations, making them the best fit for vision-oriented inference workloads where aggregate inference performance is the primary goal.

Systems with P-cores remain a capable and predictable alternative, offering solid inference performance without the throughput advantages seen in systems with AP-cores. Systems with E-cores can support the inference pipeline, but they are better aligned with workloads that emphasize efficiency and concurrency over maximum inference throughput.

NAMD

NAMD is a parallel molecular dynamics application widely used in scientific research, pharmaceutical development, and HPC environments. It relies heavily on multithreaded execution, memory bandwidth, and sustained floating-point throughput to simulate complex biological systems. NAMD scales efficiently across many cores and benefits from high aggregate compute density. For these reasons, it provides an effective benchmark for evaluating how Intel Xeon processor architectures handle large, compute-intensive workloads that stress both the CPU pipeline and the memory subsystem.

Our NAMD workload testing focused on evaluating how the core architectures of different Intel Xeon processors handle large-scale, highly parallel molecular dynamics simulations. The test emphasized sustained floating-point computation, thread scalability, memory bandwidth utilization, and inter-core communication efficiency. These are key factors that influence performance in HPC environments. For this testing, we subjected each core type to identical simulation parameters, allowing the benchmarking process to examine core scheduling behavior, Single Instruction, Multiple Data (SIMD) instruction usage, and long-duration computational stability. This approach provided a consistent framework for assessing how heterogeneous Intel Xeon processor architectures manage the technical demands of NAMD’s multithreaded simulation pipeline.

Results

Initial per-job measurements show that systems with P-cores deliver slightly higher single-job throughput per core than systems with AP-cores across the tested NAMD workloads. Viewed in isolation, individual jobs complete marginally faster on P-cores. That perspective, however, does not capture total work completed once job concurrency and aggregate throughput are taken into account.

Systems with AP-cores, with their substantially higher core counts, ran 30 simultaneous 8-core NAMD jobs, compared to 16 concurrent jobs on systems with P-cores. When total nanoseconds simulated per day are summed across all running jobs, systems with AP-cores consistently produced higher overall simulation throughput. Across molecular-simulation systems, both with and without Intel® Advanced Vector Extensions (Intel® AVX) instructions enabled, systems with AP-cores delivered between 84% and 87% more total simulation work than systems with P-cores over the same time period.

In practice, this means that while individual NAMD simulations might complete slightly sooner on systems with P-cores, systems with AP-cores deliver significantly greater total throughput when many independent simulations are run concurrently.

Conclusions

Evaluating performance based solely on per-job execution speed can be misleading for NAMD and similar HPC workloads. Although systems with P-cores offer modest advantages for single-job runs, real-world NAMD usage more often involves ensembles of independent simulations rather than a single, long-running job.

In these scenarios, Intel Xeon processors with AP-cores are a strong fit for research, engineering, and life-sciences environments where maximizing aggregate scientific throughput matters more than minimizing the runtime of any individual simulation.

Choosing the Right Processor for Your Workloads

Modern enterprise workloads call for flexibility rather than a one-size-fits-all approach to performance. Hybrid Intel Xeon processor architectures make it possible to align core capabilities with specific business needs, whether the priority is responsiveness, efficiency, or overall throughput.

E-Cores

Our testing shows that E-cores are optimized for workloads that emphasize density, efficiency, and parallel scalability rather than peak per-thread speed. In web-serving and API-style workloads, WRK results showed that systems with E-cores delivered throughput comparable to systems with P-cores and those with AP-cores across a wide range of concurrency levels, even without relying on high-performance cores.

These characteristics make E-cores a strong fit for front-end services, microservices, and other highly parallel workloads built around lightweight requests. Offloading background, asynchronous, or non-critical tasks to E-cores can reduce power consumption, improve thermal behavior, and enable more predictable infrastructure scaling.

For cloud providers, software-as-a-service (SaaS) platforms, and large enterprise IT environments, systems with E-cores offer a practical way to improve efficiency without compromising service quality. Moreover, E-cores can help organizations of all types meet sustainability goals and comply with emerging regulatory requirements related to power consumption and carbon emissions.

P-Cores

P-cores perform best in environments where strong single-thread performance and predictable execution are critical. In compute-intensive workloads such as NAMD, systems with P-cores delivered slightly higher per-job performance and exhibited stable, interpretable behavior across a broad range of configurations.

P-cores are particularly well-suited for latency-sensitive applications that depend on fast response times and consistent per-task execution, including real-time analytics, transactional systems, and performance-tiered services.

For organizations and business environments that differentiate offerings based on responsiveness or per-request latency, systems with P-cores provide a dependable foundation without requiring the operational complexity associated with large-scale parallel execution. P-cores:

- Support premium product tiers and flagship applications that demand top-end responsiveness.
- Enable new business scenarios such as real-time field service analytics, retail intelligence platforms, and industrial automation workloads.

AP-Cores

AP-cores are designed for workloads that benefit from high core counts and sustained multithreaded throughput. In our testing, systems with AP-cores consistently delivered the highest aggregate performance for CPU-based AI inference using the OpenVINO toolkit with ResNet-50, outperforming both systems with P-cores and those with E-cores across configurations. Systems with AP-cores also achieved substantially higher total simulation throughput for NAMD by running many independent jobs in parallel.

This combination makes AP-cores well-suited for vision-style inference, HPC-class simulations, complex analytics, and research workloads where total work completed per unit of time matters more than single-job speed. By consolidating compute capacity, memory bandwidth, and accelerator-class capabilities within a single platform, systems with AP-cores can reduce hardware sprawl and simplify deployment without incurring the power and thermal overhead typically associated with discrete accelerators.

In business environments, systems with AP-cores are a strong choice for research, engineering, and data-science teams that need to maximize aggregate throughput while scaling efficiently across large problem sets. AP-cores:

- Consolidate multiple discrete components into a unified design, lowering hardware spend and simplifying system architecture.
- Reduce integration complexity and accelerate deployment of advanced research, simulation, and enterprise AI workloads.

Explore

To further explore Intel Xeon processors, visit **www.intel.com/xeon**.

Refer to the **Methodology** document for a summary of the steps taken for our comparisons.

FAQ

1. Who is this report intended for?

This report is intended for IT leaders, infrastructure architects, and technical decision-makers evaluating server platforms for enterprise environments. It helps readers understand how Intel Xeon processors with E-cores, P-cores, and AP-cores enhance workload performance, efficiency, and scalability to support informed infrastructure and workload-placement decisions.

2. What are the key differences between E-cores, P-cores, and AP-cores in real-world workload performance?

E-cores provide exceptional performance per watt and are optimized for scalable, parallel, and background compute tasks. P-cores deliver high single-thread performance and excel at latency-sensitive workloads such as AI inference, real-time analytics, and interactive applications. AP-cores are designed for large multithreaded workloads and integrate accelerators and high-bandwidth memory to support simulations, AI pipelines, and HPC-style analytics.

3. How does workload-specific core assignment improve performance, efficiency, and reliability in enterprise environments?

Assigning each workload to the most suitable core type reduces resource contention, improves overall throughput, and lowers energy consumption. This approach ensures E-cores handle scalable or background operations, P-cores are reserved for critical low-latency tasks, and AP-cores accelerate heavy compute pipelines—resulting in more predictable performance and better infrastructure utilization.

4. Which types of workloads benefit most from AP-cores, and why?

AP-cores are ideal for simulation-heavy, accelerator-enhanced, and memory-intensive workloads. Applications such as molecular dynamics (NAMD), large-scale analytics, AI model tuning, and scientific or engineering simulations can benefit from their high multithreaded throughput, consistent performance under sustained load, and efficient integration of NPUs, GPUs, and high-bandwidth memory.

5. What types of workloads were tested in this study, and why were they chosen?

The study tested three representative workloads: WRK for web server and API performance, the OpenVINO toolkit with ResNet-50 for AI inference, and NAMD for HPC simulation pipelines. These workloads were selected because they reflect real-world enterprise use cases across web services, AI/machine learning (ML), and HPC, allowing for a practical evaluation of how different Intel Xeon processor core types (E-cores, P-cores, and AP-cores) handle specific computational demands.

6. How does understanding core-specific performance in these tests help businesses optimize their infrastructures?

By analyzing how E-cores, P-cores, and AP-cores perform under these representative workloads, IT teams can make informed decisions about workload placement, maximizing throughput, minimizing latency, and improving performance per watt. This insight enables organizations to optimize server utilization, reduce energy costs, and ensure that compute-intensive tasks are assigned to the most appropriate cores, ultimately improving operational efficiency and user experience.

Appendix: Testing Setups

Model	Dell PowerEdge R770 (with E-Cores)	Dell PowerEdge R770 (with P-Cores)	Dell PowerEdge R770AP (with AP-Cores)
CPU	2 x Intel Xeon 6740E processors (96 cores each), 250 W thermal design power (TDP)	2 x Intel Xeon 6767P processors (64 cores each), 350 W TDP	2 x Intel Xeon 6978P processors (120 cores each), 500 W TDP
Memory	2,048 GB (32 x 64 GB) DDR5 RDIMM @ 5,200 MT/s (2DPC)	2,048 GB (32 x 64 GB) DDR5 RDIMM @ 5,200 MT/s (2DPC)	1,536 GB (24 x 64 GB) DDR5 RDIMM @ 6,400 MT/s (1DPC)
Power	PSU 1: 1,500 W PSU 2: 1,500 W	PSU 1: 1,500 W PSU 2: 1,500 W	PSU 1: 2,400 W PSU 2: 2,400 W
Networking	4-port, 25 Gb Broadcom® SFP 57504S OCP network interface controller (NIC)	4-port, 25 Gb Broadcom SFP 57504S OCP NIC	2-port, 100 Gb Mellanox® ConnectX®-6 Dx NIC Broadcom BCM57608 2x200 Gb OCP Ethernet NIC—8C:84:74:9D:89:16
Storage	2 x 480 GB Micron EC NVM Express® (NVMe®) ISE 7450 RI M.2 80 (Dell™ Boot-Optimized Server Storage [BOSS] attach) 8 x 3.84 TB Micron MTFDDAK3T8TGA-1B (Dell™ PowerEdge RAID Controller [PERC] H365i attach) 4 x 1.92 TB KIOXIA® Corporation DC NVMe CD8 U.2 (CPU direct attach)	2 x 480 GB Micron EC NVMe ISE 7450 RI M.2 80 (Dell BOSS attach) 8 x 3.84 TB Micron MTFDDAK3T8TGA-1B (Dell PERC H365i attach) 4 x 1.92 TB KIOXIA DC NVMe CD8 U.2 (CPU direct attach)	2 x 480 GB Micron EC NVMe ISE 7450 RI M.2 80 (Dell BOSS attach) 2 x 1.6 TB Samsung® DC NVMe PM9D5a MU U.2 (CPU direct attach)

Endnotes

[1]NAMD was developed by the Theoretical and Computational Biophysics Group in the Beckman Institute for Advanced Science and Technology at the University of Illinois at Urbana-Champaign. **www.ks.uiuc.edu/Research/namd/**.

[2]Where parity is defined as being within ±3% performance per watt.