Research Abstract

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

Research by Prowess Consulting shows measurable benefits when IT teams target workloads across different Intel Xeon processor types.

IT teams in modern data centers must support a mix of high-concurrency services, latency-sensitive applications, and increasingly compute-intensive AI and simulation pipelines, often within fixed power, cooling, and space envelopes. Intel Xeon processors address these competing demands with three specialized core types: Efficient-cores (E-cores), Performance-cores (P-cores), and Advanced Performance–cores (AP-cores). Each core type is optimized for different execution patterns and operational priorities.

In this study, Prowess Consulting evaluated how systems based on Intel Xeon processors with E-cores, P-cores, and AP-cores behave across three representative workloads:

- WRK® web benchmark (WRK2) for web concurrency
- ResNet®-50 using Intel®-optimized software (including the Intel® Distribution of OpenVINO™ toolkit) for CPU-based AI inference
- Nanoscale Molecular Dynamics (NAMD) for high-performance computing (HPC)-style simulation throughput[1]

The goal of our testing was to determine how performance and efficiency vary when workloads are intentionally matched to the core types best suited to their execution profiles.

Highlights

We found the following workload impacts across E-cores, P-cores, and AP-cores:[2]

WEB-SERVING THROUGHPUT (WRK2): Systems with E-cores achieved throughput within **0.25%** of systems with P-cores (with up to 10,000 requests at 50 connections), indicating that concurrency-heavy web workloads can deliver strong throughput without requiring performance-focused cores.

RESNET-50: Systems with AP-cores delivered up to **120% higher** CPU-based AI inference throughput than systems with P-cores across tested precision modes and block sizes.

NAMD: Systems with AP-cores completed **84%–87%** more molecular simulation work than systems with P-cores by running more jobs concurrently (30 simultaneous 8-core jobs versus 16).

EFFICIENCY NUANCE (RESNET-50): Although systems with AP-cores led in raw inference throughput, **44%** of tested precision/block-size combinations showed **systems with P-cores at parity with or exceeding systems with AP-cores in performance per watt,** which emphasizes the importance of considering both throughput and total platform power.

Overall, these results reinforce that getting the best outcome depends on the workload objective. E-cores are a strong fit for highly parallel services, where density and efficiency matter most. P-cores provide responsive, predictable performance for latency-sensitive use cases. AP-cores excel when organizations need maximum aggregate throughput for CPU-based AI inference and large-scale multithreaded simulations, where total work completed matters more than the runtime of any individual job.

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]For details, see the full study: Prowess Consulting. "**Benefits and Comparison of Intel® Xeon® Processors with E-Cores, P-Cores, and AP-Cores for Select Workloads**." April 2026.

0426/250111