Hewlett Packard Enterprise | intel

Brochure

HIGH PERFORMANCE COMPUTING FOR MANUFACTURING

HPE and Intel's partnership empowers your infrastructure, optimization, and consumption models

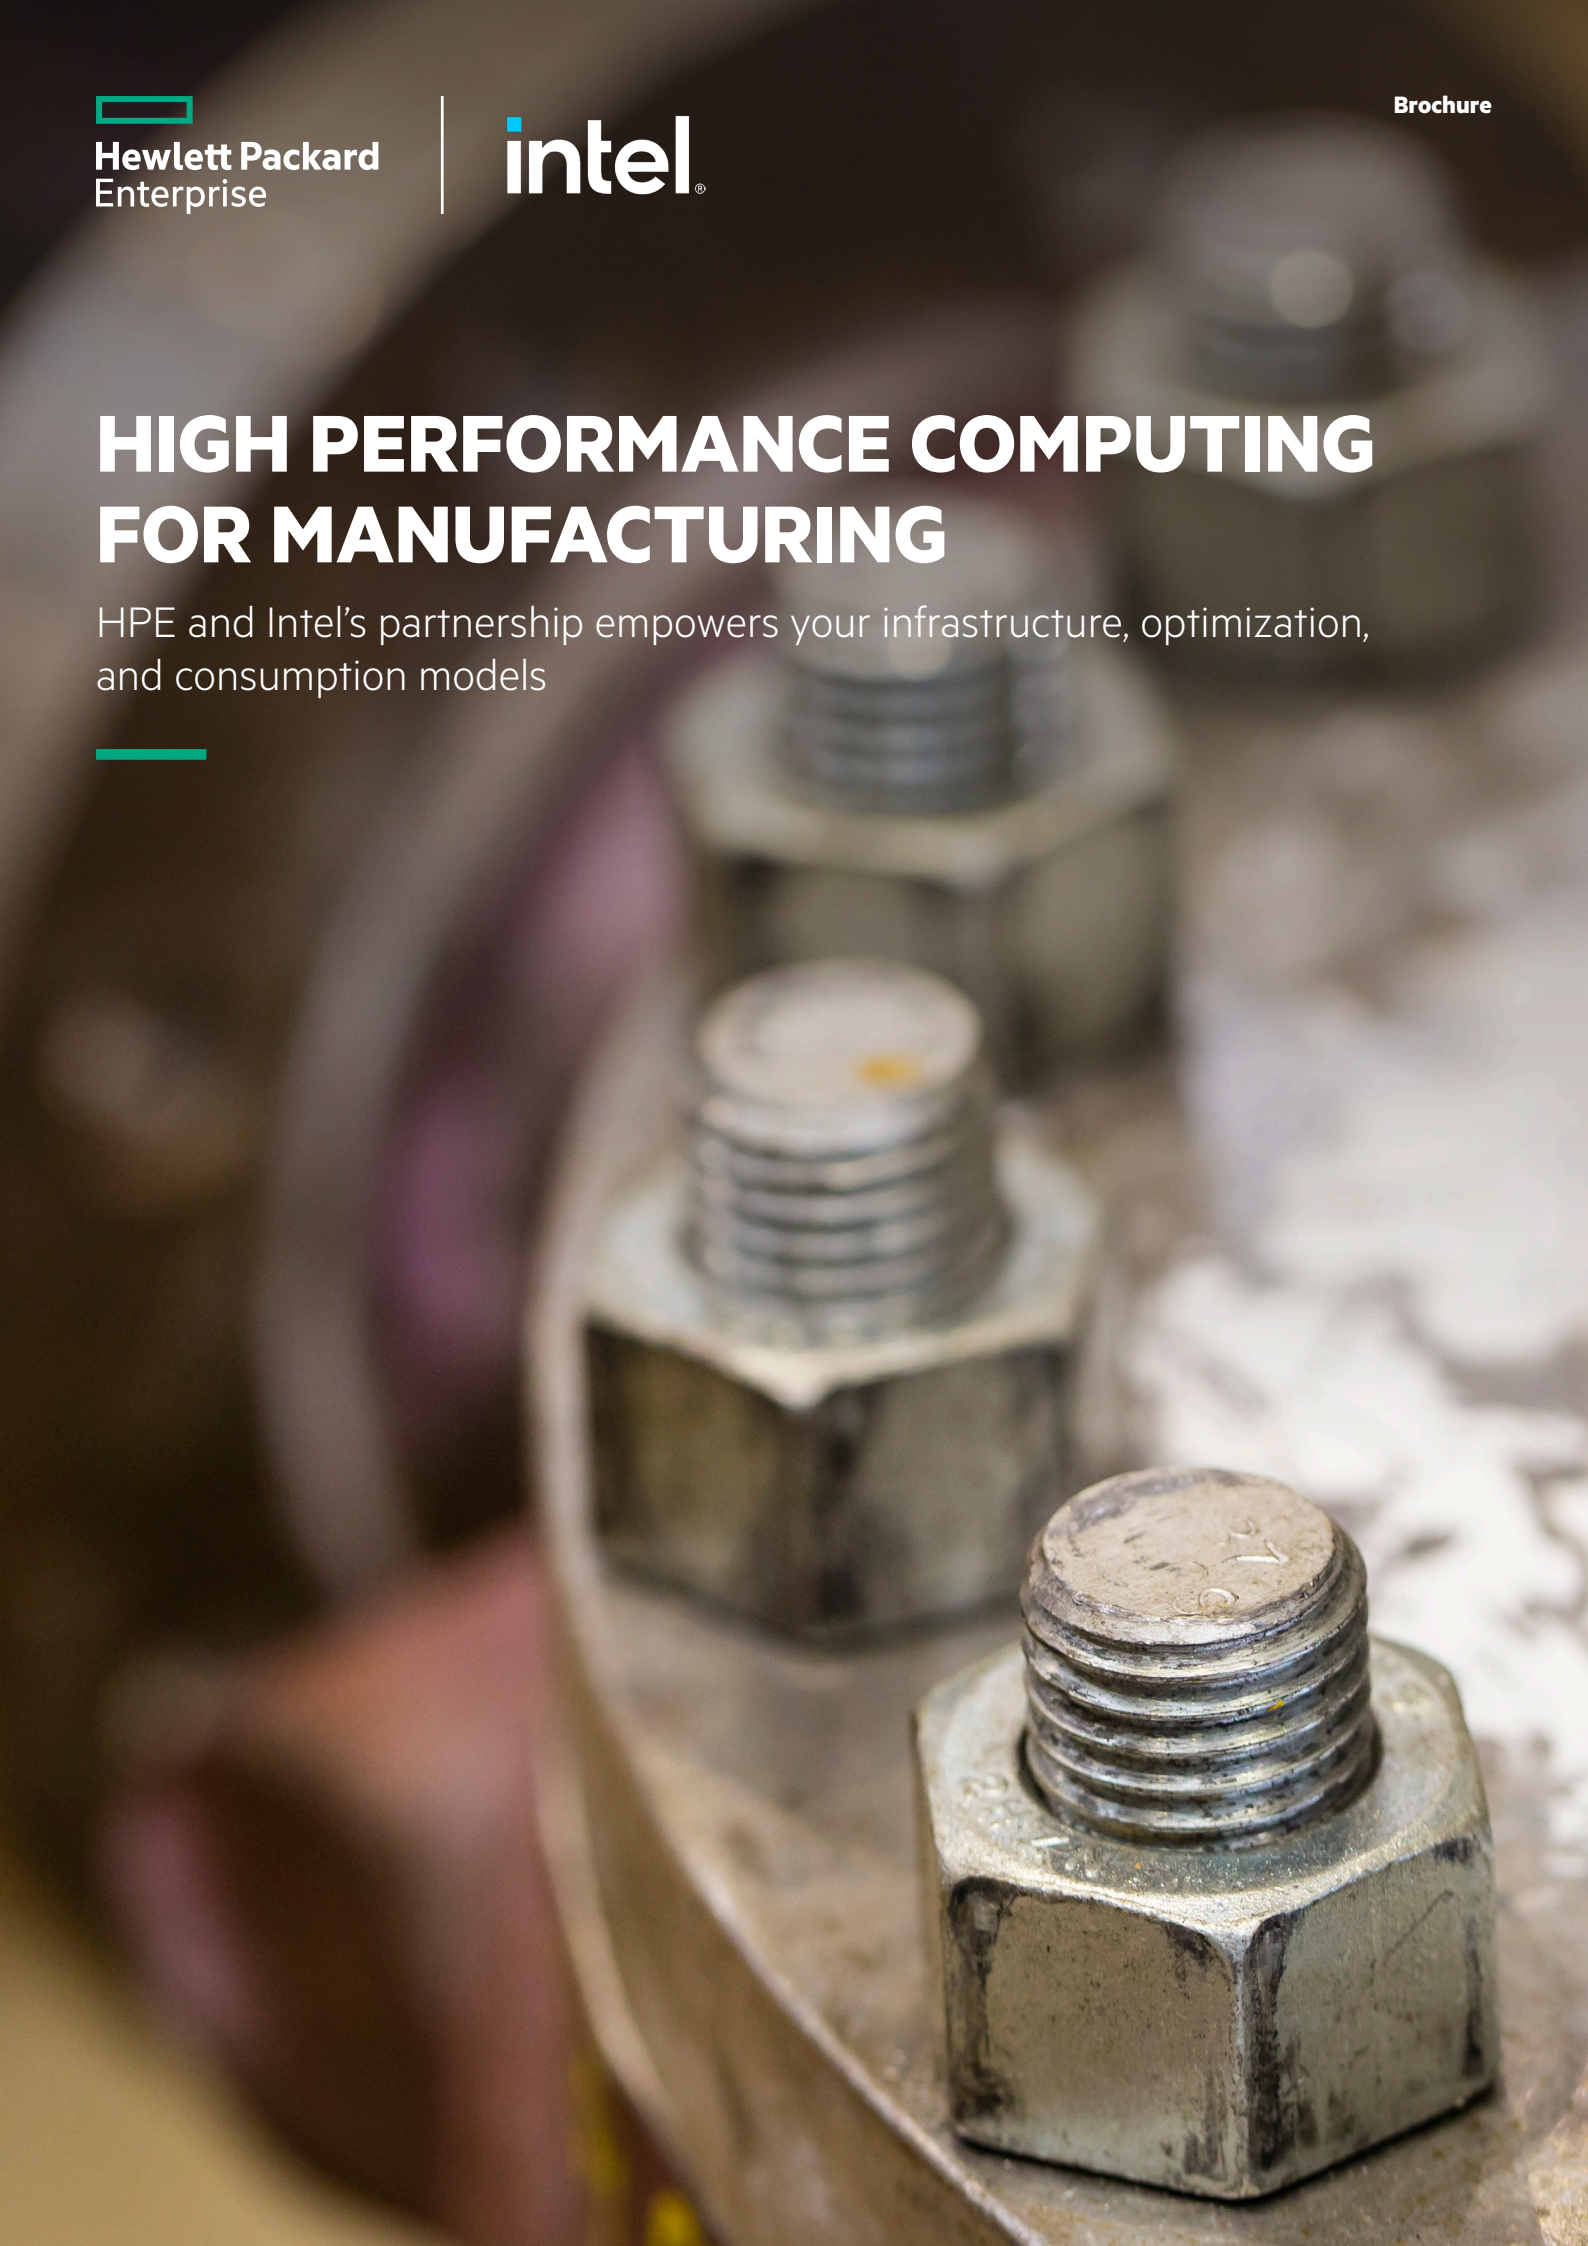

Check if the document is available in the language of your choice.

WHEN THE ROUTE TO MARKET BECOMES A RACE

Bringing a product to market today requires an increasingly fine balance of speed, quality, and cost efficiency. Engineering and development cycles must be squeezed to reduce time to market, but productivity cannot be compromised since customers expect higher quality and more choice than ever before. Meanwhile competitive pricing pressure means margins must be protected at all costs, so design teams must work as efficiently as possible.

The increasingly vast volume of data available to and generated by manufacturers running computer aided engineering (CAE) simulations can be the key to successfully balancing these imperatives. Being able to manage the data in more detail and at greater speed, using tools like artificial intelligence (AI), means engineers can make design choices faster and with greater confidence. It also builds in flexibility to respond quickly to changes in demand.

As data volumes continue to grow, and simulations become higher resolution and ever-more complex, high performance computing (HPC) becomes increasingly important in engineering efficiency and quality. With HPC, designers can move from component- or part-level simulations to large-scale multiple physics and full body multiphysics simulations using a range of disciplines. These include:

- Computational fluid dynamics (CFD)
- Computational structural mechanics (CSM) for both implicit and explicit finite element analysis
- Computational electromagnetics (CEM) solvers from major CAE applications such as Altair Radioss and Altair Feko; Ansys Fluent, Ansys Mechanical, Ansys LS-DYNA and Ansys HFSS; MSC Nastran; Siemens STAR-CCM+; OpenFOAM™ and SIMULIA® Abaqus FEA®; and ESI PAM-CRASH (VPS).

HPE AND INTEL COLLABORATE TO SUPPORT HPC

The strategic alliance of HPE and Intel® is embodied in our long-term commitment to HPC and to companies who rely on it for business success.

HPE and Intel innovate across processing, storage, fabric, software and workload optimization and consumption models. This makes HPC infrastructure easier to deploy, manage, and scale. As a result, HPC users can leverage their hardware and software investments to reach valuable insights faster.

Intel and HPE have a 30-year history of collaborating to develop tried and tested solutions. Their forward trajectory is well aligned, with HPE focusing on offering 'everything as a service' by 2022, and ever-increasing performance and value supported by current and future generations of Intel technologies. This long-term HPC focus will help manufacturers accelerate their digital transformation as they ramp up activity to take advantage of the post-pandemic economic recovery.

An effective way of running CAE workloads like CFD, CSM and CEM is using CAE ISV software optimized on Intel architectures, so an Intel-based HPE HPC platform provides an ideal infrastructure.

BETTER TOGETHER—INTEL AND HPE PARTNER DELIVER VALUE TO HPC MANUFACTURING CUSTOMERS

HPE—Delivering innovation on-demand at any size

HPE systems based on Intel products are engineered to optimize performance for supercomputing systems of any size for every data center, with systems that perform like a supercomputer and run like a cloud—giving you the power and flexibility to run increasingly complex models and more refined simulations quicker than ever.

HPE leverages technology from Intel—including 3rd Gen Intel® Xeon® Scalable processors—to deliver supercomputing capability at any scale and footprint—from cloud to data center.

- HPC as a Service: Run your most demanding workloads with fully managed, pre-bundled HPC cloud services to operate in any data center or co-location environment.

- HPE Apollo 20 System: Purpose-built, high density platform delivering advanced performance for compute-intensive HPC and AI workloads.
- HPE Apollo 2000 System: Density optimized, scale-out compute platform designed to speed CAE workloads and bring new products and services to market faster.

HPE and Intel also bring a suite of enabling technologies in interconnect and storage which create a fully tuned system.

INTEL—SCALABLE, OPTIMIZED PERFORMANCE FOR MANUFACTURING HPC

Intel's comprehensive HPC offering with end-to-end capabilities to address a broad range of customer needs. The diverse architectures within our portfolio support varying workload types and requirements, letting you address increasingly large and complex data sets quickly. This means you can get to market faster and more cost-effectively, run more iterations on complex and higher-res models to get faster results, and run diverse workloads on the same HPE cluster, reducing time on system and cost.

HPE AND INTEL POWER MANUFACTURING HPC IN SIX KEY AREAS

Processing power and flexibility
With built-in HPC and AI acceleration, 3rd Gen Intel Xeon Scalable processors deliver up to 1.52x higher performance for manufacturing workloads than previous generations,[1] helping speed time to results and optimize per-core software licensing costs. Built-in Intel Speed Select technology provides three configurations in one CPU to address a broader range of workload demands.

Security
Intel Software Guard Extensions (Intel SGX) provides silicon-level protection for sensitive/IP data, delivering a security foundation for collaboration and federated learning in manufacturing.

Optimization
Many of the most important applications in design engineering are optimized for 3rd Gen Intel Xeon Scalable processors.

Scale
Purpose-built, density-optimized HPC platforms with HPE Performance Cluster Manager enable straightforward, affordable scaling to meet your growing demands.

Storage
Intel® Optane™ persistent memory on HPE systems brings massive data sets closer to the CPU for faster time to insight. It increases data center resiliency and significantly accelerates database restart times.

Consumption model
HPE GreenLake brings the cloud experience—self-serve, pay-per-use, scale up and down, and managed for you by HPE and our partners—to applications and data everywhere, in edges, colocations and data centers. This enables you to free up capital, boost operational and financial flexibility and free up talent.

[1] Geomean across geomean Altair Radioss, Ansys Fluent, Ansys LS-DYNA, Converge, NUMECA, and OpenFOAM workloads. See [105] at intel.com/3gen-xeon-config for details. Results may vary.

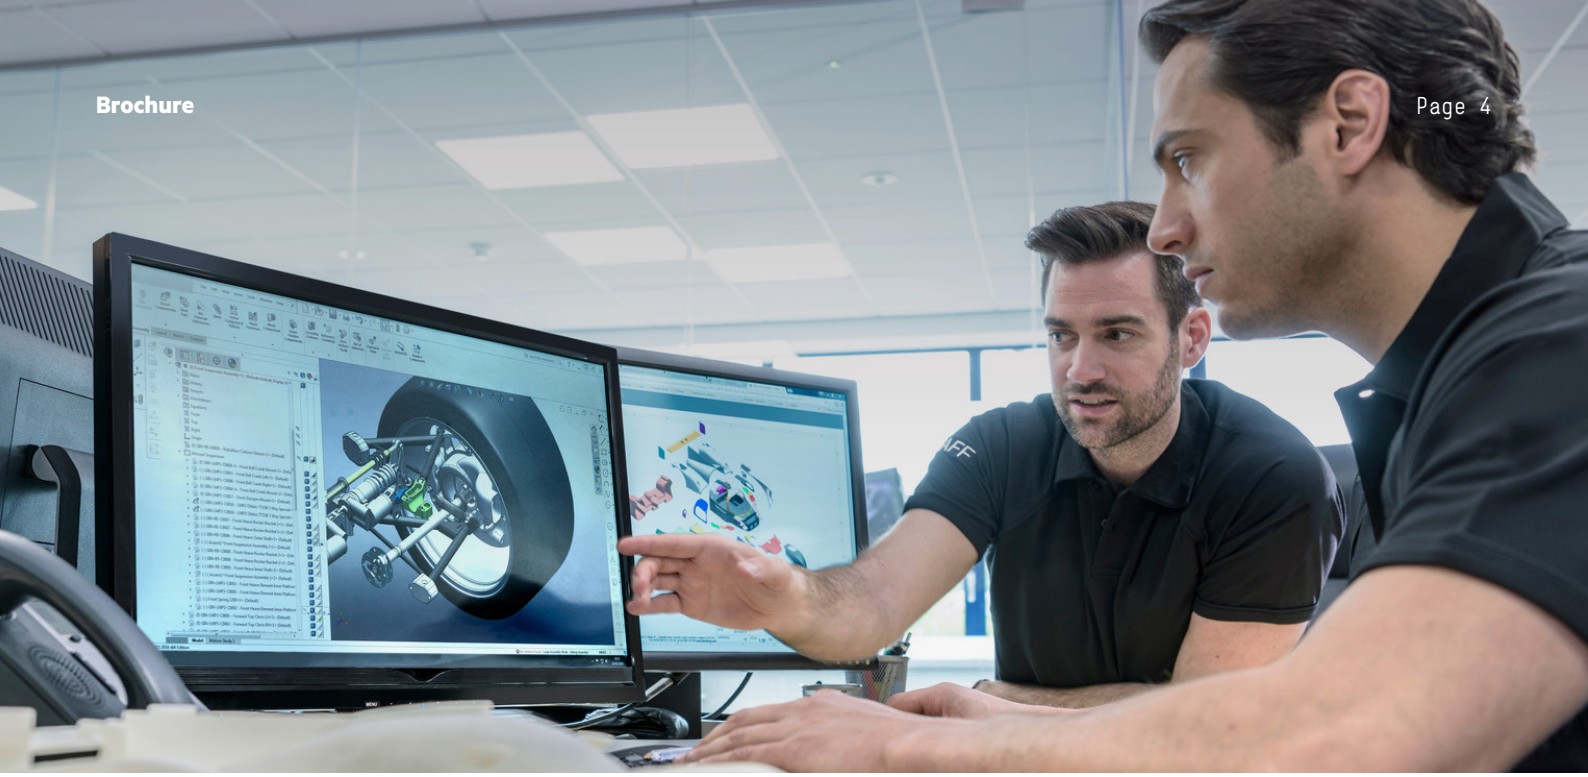

NEXT-GEN CFD POWERED BY HPE AND INTEL-BASED HPCAAS[2]

Advania Data Centers specializes in HPC-as-a-Service, helping customers perform complex, data-driven scientific, design and engineering projects. Advania uses HPE Apollo server hardware and Intel Select Solutions for HPC & AI Converged Clusters (Magpie), which feature Intel Xeon Scalable processors.

NUMECA International will run its next-gen CFD software on Advania's HPCaaS platform. It chose Advania's managed HPC solutions for several benefits, including scalability, end-to-end security, and the ability to configure an isolated cluster down to bare metal. The collaboration will deliver an end-to-end CFD on HPC solution that offers NUMECA's end users a platform for their most challenging simulation and modelling projects.

Intel's end-to-end range of technologies and optimizations includes:

Compute to meet your exact requirements, and speed up AI
3rd Gen Intel Xeon Scalable processors deliver outstanding performance across a wide range of CAE workload requirements, providing an excellent foundation for both dedicated and multipurpose systems. They offer up to 40 cores and eight memory channels, including a wide range of frequencies, addressing broad power-consumption requirements. Built-in acceleration and flexibility allow Intel-based systems to flex to meet varying performance requirements, so HPC and AI workloads run optimally.

3rd Gen Intel Xeon Scalable processors are the only data center CPU with built-in AI acceleration, supported by end-to-end data science tools and a vast ecosystem of smart solutions. They deliver 1.5x more performance than other CPUs across 20 key customer AI workloads.[3] The Intel Xeon platform is engineered to deliver better performance per core, optimizing license usage for some of the most important and widely used CAE apps. And with Speed Select Technology, you get three configurations in one CPU so your infrastructure flexes to support more workloads.

[2] intel.com/content/www/us/en/customer-spotlight/stories/numeca-customer-story.html

[3] See [43] at intel.com/3gen-xeon-config for details. Results may vary.

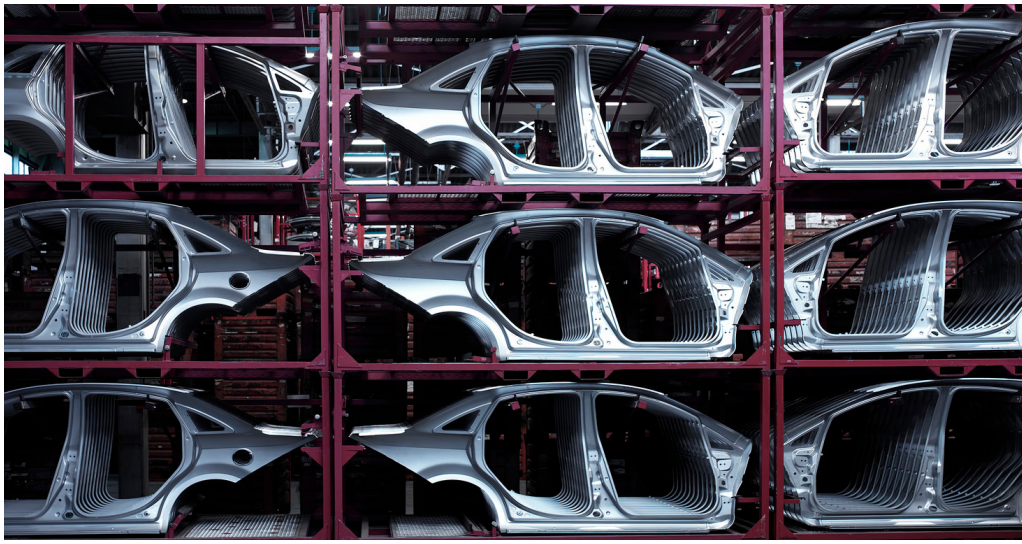

Silicon-level security
The Intel Xeon Scalable platform includes Intel Software Guard Extensions (Intel SGX) to help protect data and application code in real time from the edge to the data center and multi-tenant public cloud, enabling enhanced collaboration (for example for federated learning in AI) using shared data more securely. This means it is easier for engineers and stakeholders across the value chain to collaborate and share data or AI models while enhancing privacy and security.

Further enhancing security, Intel Crypto Acceleration increases the performance of encryption-intensive workloads including SSL web serving, 5G infrastructure, VPN/firewalls and reduces the performance impact of pervasive encryption. So business-critical intellectual property (IP) can be kept secure and still processed lightning fast.

Fast storage
The Intel Optane SSD P4800X is the world's fastest SSD for the data center.[4] With unprecedented endurance coupled with PCIe Gen 4 support, it is ideal for hot-tier caching and accelerating all types of storage solutions. Intel Optane persistent memory (PMem) 200 series is Intel's next-generation persistent memory module that provides large capacity and native persistence to help extract more value from larger datasets. The technology increases speed and agility by bringing massive data sets closer to the CPU. It increases data center resiliency and takes database restart times from hours down to minutes or seconds. This allows access to use cases that were previously unavailable with traditional storage. Data can now be stored locally and accessed remotely to provide fast, high-capacity, immediate, consistent in-memory databases.

The 3rd Gen Xeon Scalable platform supports the latest Intel Optane PMem 200 series, delivering up to 6 TB total memory per socket and on average 32 percent higher memory bandwidth compared to previous generation.[5]

[4] See [14] at edc.intel.com/content/www/us/en/products/performance/benchmarks/intel-optane-ssd-p5800x-series/?r=94788905 for details. Results may vary.

[5] edc.intel.com/content/www/us/en/products/performance/benchmarks/intel-optane-persistent-memory-200-series/.

Multiple Intel-optimized CAE applications

Many of the applications that engineers and product designers rely on are optimized for and accelerated by 3rd Gen Intel Xeon Scalable processors. Intel AVX-512 and other code optimizations can be used to help enhance performance for key applications like Altair Radioss, MS Nastran, Abaqus/Standard, and Ansys Mechanical. For example, Ansys and Intel engineers optimized Fluent to leverage AVX-512 instructions, resulting in performance gains of up to 19 percent.[6] Developers working with these packages have access to a wide range of Intel compilers and libraries.

Contact your HPE representative to learn more.

LEARN MORE AT

hpe.com/info/hpc-manufacturing

[6] intel.com/content/www/us/en/high-performance-computing/ansys-fluent-simulations-at-top-speed-paper.html.

Make the right purchase decision.
Contact our presales specialists.

Get updates

Hewlett Packard Enterprise

a50004594ENW, July 2021