Brochure

Tackle data-intensive HPC problems holistically

HPE Compute Scale-up Server 3200 and
HPE Superdome Flex family for HPC use cases

Hewlett Packard Enterprise | intel

Accelerate discovery

Organizations across the globe utilize high performance computing (HPC) to solve difficult problems in science, engineering, and business. Many depend on the robust HPE ecosystem of HPC solutions to gain the computational power they need and to accelerate time to discovery. As HPC problems grow in size and diversity, HPE solutions have evolved to address the full continuum of data challenges that spans high-velocity data processing, real-time analysis, and insights that drive progress.

HPC workloads are commonly run on clustered systems, in which computational problems can be distributed across multiple servers (nodes) working in parallel and which are connected over a high-speed network with shared storage. In addition to extending computational capabilities by orders of magnitude, HPC teams can run many types and sizes of jobs within and across nodes concurrently.

Complex HPC problems, especially those that are data-intensive and feature strong data interdependence, are challenging to run across multiple nodes. Moving data between nodes in a scale-out cluster is slow; latency is an issue, and it takes time to re-assemble large and complex datasets that were broken up to enable distribution. These problems are better tackled by a single "shared-memory" node, also known as a "fat" node, which does not need to break jobs into pieces across a cluster. Examples of such workloads include:

Computer-aided engineering (CAE) such as optimizing placement of airplane antennae and GPS devices to minimize interference using electromagnetic simulation, improving the structural design of wind turbines using FEM[1] simulation, designing screw augers for transporting bulk material using DEM[2] technology, or perfecting race car designs using computational fluid dynamics to simulate aerodynamics.

Medical and agricultural genomics such as genome mapping, a process that compares billions of small sequences and terabytes of data with a previously assembled genome until complete, genomic research in which a full DNA sequence is analyzed to identify gene functionality and variations to help predict, diagnose, and treat disease, or precision medicine, an approach that considers individual gene variabilities to provide personalized patient treatment.

Cybersecurity, such as detection of lateral movement attacks, which challenges teams to traverse and identify threats that often look like normal network traffic. In a lateral movement attack, a bad actor gains access to and moves through the network searching for vulnerabilities. To reduce Mean Time To Detection (MTTD), teams must be able to scan cyber logs more quickly.

Fraud prevention, such as scanning and comparing barcodes with previously scanned barcodes (thousands to billions) to identify copies before loss or damage is incurred.

In situ visualization, in which a simulation is run and the data generated is analyzed visually and in real time.

[1] Finite Element Method

[2] Discrete Element Method

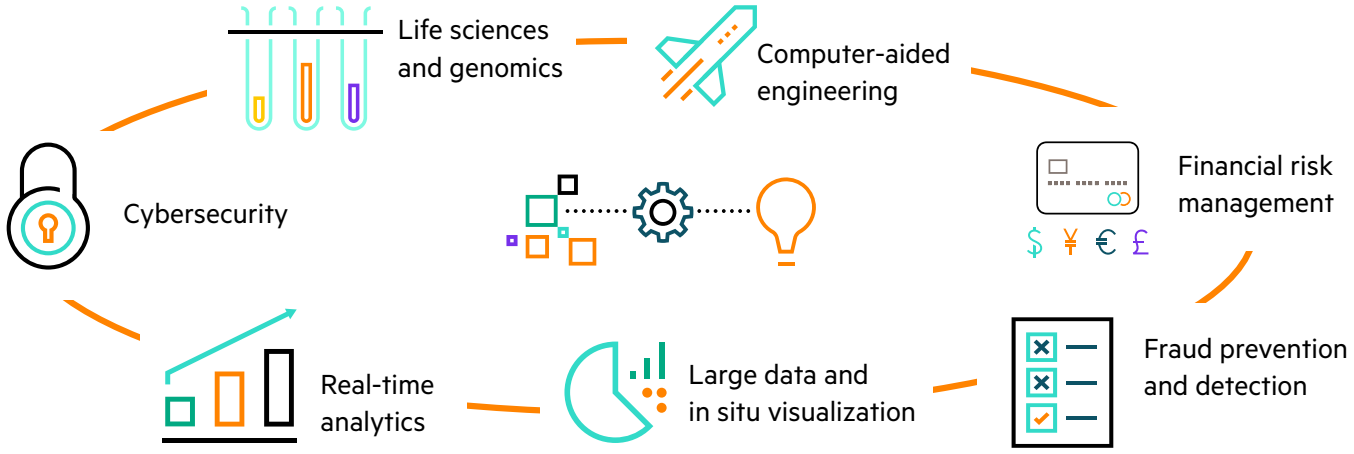

Figure 1. A sample of data-intensive HPC use cases

Workload challenge

Accommodating such workloads on the largest nodes can be challenging for HPC teams as they often compete for time, leaving some jobs waiting and less-powerful nodes underutilized. This increases time and effort, and reduces efficiency.

Such workloads might also be too big for the node to handle. If the node's memory is exhausted, the job will fail, wasting hours of running time. To fit on the node, models, and simulations must then be made smaller, for example by using sections and estimations or decreasing granularity. This, however, reduces accuracy, increases prototype costs, and adds time to production or discovery.

Staff challenge

An additional challenge for relatively small HPC teams is the question of who will manage the cluster environment, a role that includes balancing workloads, tuning performance, and other IT tasks. Can the IT department dedicate a cluster administrator, or must an engineer or scientist take time from development and research to learn and perform?

Solution: shared-memory HPC with HPE scale-up systems powered by Intel®

The ideal solution for solving complex, data-intensive problems holistically is with shared-memory HPC and the breakthrough SMP (symmetric multiprocessing) of the HPE Compute Scale-up Server 3200 and HPE Superdome Flex family, powered by Intel® Xeon® Scalable processors. Providing unparalleled scale-up compute and shared memory resources with single-system simplicity, these platforms equip your HPC team to:

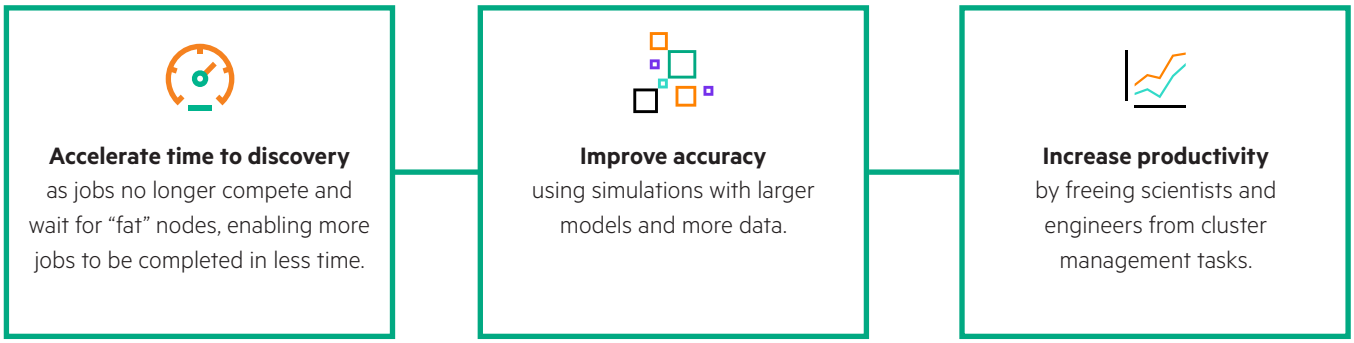

HPE Compute Scale-up Server 3200

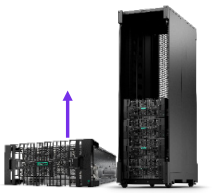

- For midsize to large enterprise environments
- 4–16 sockets as a single system
- Scale in 4-socket increments
- 4th Gen Intel Xeon Scalable processors
- Designed to provide 128 GB to 32 TB of shared memory
- Up to 64 PCIe 5.0 cards

HPE Superdome Flex 280

- For small to midsize enterprise environments
- 2–8 sockets as a single system
- Scale in 2-socket increments
- 3rd Gen Intel Xeon Scalable processors
- Designed to provide 64 GB to 24 TB of shared memory
- Up to 32 PCIe 3.0 cards

HPE Superdome Flex

- For midsize to large enterprise environments
- 4–32 sockets as a single system
- Scale in 4-socket increments
- 2nd Gen Intel Xeon Scalable processors
- Designed to provide 768 GB to 48 TB shared memory
- Up to 128 PCIe cards

Figure 2. HPE Compute Scale-up Server 3200 and HPE Superdome Flex family

Software developers will also greatly benefit from HPE shared-memory HPC solutions because, unlike when moving from a desktop to a cluster, there is no need to modify code. These HPE platforms are like giant Linux® workstations with lightning speed.

These HPE scale-up servers powered by Intel deliver:

- **Scale from small to huge.** A unique modular architecture allows you to start small and grow at your own pace without sacrificing performance. HPE Compute Scale-up Server 3200 can scale seamlessly from 4 to 16 sockets in 4-socket increments. Its shared memory capacity ranges from 128 GB to 32 TB and it utilizes ultra-fast DDR5 DIMMs. HPE Superdome Flex 280 offers a low entry point and granular scaling, starting at 2 and scaling to 8 sockets in 2-socket increments. HPE Superdome Flex scales from 4 to 32 sockets in 4-socket increments.
- **High performance at scale.** Achieve the performance levels required by complex HPC problems with ultra-low latency and high bandwidth technology plus the power of Intel Xeon Scalable processors.
- **Optimum flexibility.** The family's modularity and scale help you avoid overprovisioning, disruptive upgrades, and the cost and complexity those carry.
- **Unbounded I/O.** Capitalize on a well-balanced I/O subsystem for high performance. HPE Compute Scale-up Server 3200 offers I/O flexibility with three options to support a variety of workloads. HPE Superdome Flex 280 and HPE Superdome Flex, offer 16 or 12 slots per 4-socket chassis. The servers support a multitude of cards, providing rich solution flexibility.
- **Extreme availability.** The family's unique RAS features include **Firmware First**, which contains errors at the firmware level and insulates the OS; a built-in **Analysis Engine** to monitor, analyze, and fix errors with self-repair; and **advanced resiliency** across every subsystem. For maximum application availability you can add HPE's leading clustering solution **HPE Serviceguard for Linux**.
- **Trusted security.** To protect valuable data and workloads, the security strategy of these mission-critical servers focuses on helping minimize threat exposure to vulnerabilities, including those found in common firmware. The unique server management architecture provides many security advantages, for example, the operating system, which is a major source of vulnerabilities, is rarely trusted. HPE Compute Scale-up Server 3200 features silicon root of trust from HPE and adds dual Trusted Platform Module 2.0 (TPM 2.0) chips, bundled and authenticated firmware updates, and new automatic deconfiguration capabilities for the I/O subsystem.
- **Simplified user experience.** Leverage HPE OneView Management, OpenStack cluster management, industry-standard Redfish APIs, and a simplified management GUI. You can also consume HPE Compute Scale-up Server 3200 and HPE Superdome Flex family as a service, via HPE GreenLake.

Real-world results

Organizations from a wide range of industries are leveraging these HPE scale-up systems to tackle data-intensive HPC workloads holistically.

Fighting pandemics with genomics

A large US university is helping the world combat viruses and bacteria by unlocking their genomes. By adopting Intel-powered HPE scale-up platforms it achieved:

- 10x faster sequencing of virus genomes and identification of variants.
- Faster decisions around vaccine distribution and public health management.

Powering up international research

A European regional data innovation center is using HPC to speed time to discovery and increases research accuracy. Its HPE solutions on Intel provided:

- Native Kubernetes support and self-service applications.
- A single data management framework to enable collaborative, optimal experiences.

Building a genomic supercomputer

A life science and genomic medicine research center in Japan provides researchers with opportunities for joint use of and research with supercomputers. The center adopted Intel-based HPE scale-up servers that delivered:

- Huge, shared memory for connecting a large number of genome fragments and restoring sequences.
- Capability to meet researchers' needs more quickly and efficiently.

Go further, faster, with HPE high performance computing solutions

HPE is the global leader in high performance computing solutions, with deep expertise across HPC workloads and a powerful, purpose-built product portfolio that make supercomputing more accessible and affordable for organizations of all sizes.

Ask your Hewlett Packard Enterprise sales representative about equipping your team with shared-memory HPC solutions powered by Intel-based HPE Compute Scale-up Server 3200 and HPE Superdome Flex family.

Learn more at

HPE.com/superdome

Chat now (sales)

a50000523ENW, Rev. 4

Hewlett Packard Enterprise