Hewlett Packard Enterprise

Brochure

HPE SHARED MEMORY HIGH PERFORMANCE COMPUTING

Tackle complex, data-intensive HPC problems holistically with the unique scale-up architecture of HPE Superdome Flex

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

TABLE OF CONTENTS

ACCELERATE DISCOVERY

Organizations across the globe utilize high-performance computing (HPC) to solve difficult problems in science, engineering, and business. Many depend on the HPE ecosystem of accessible, affordable HPC solutions to gain the computational power they need and to accelerate time to discovery. And as HPC problems grow in size and diversity, HPE solutions equip research, development, and security teams to meet current and future requirements.

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.

Some HPC problems, however, are challenging to run across multiple nodes. Complex and often data-intensive with strong interdependence, these problems are often best tackled holistically using a single node. And not just any node, but a "fat" node with the greatest amount of processing power and memory. 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.
- **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

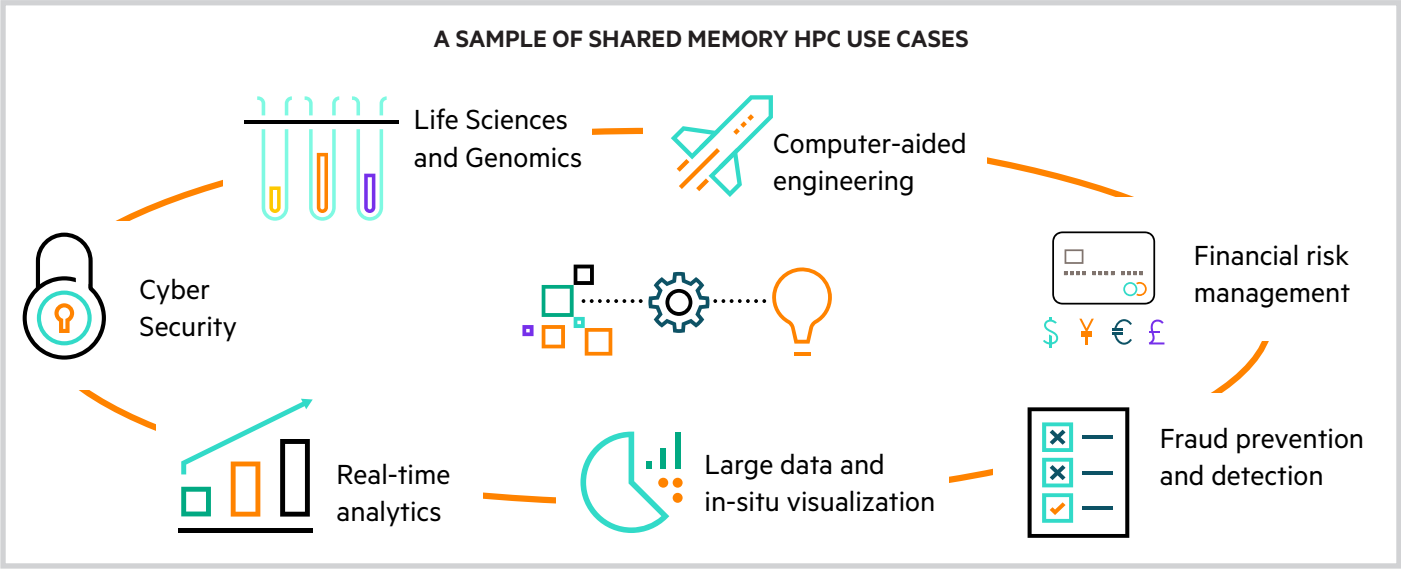

Workload challenge
Accommodating such workloads on the largest nodes can be challenging for HPC teams as they often compete for time, leaving some jobs waiting and less-powerful nodes under-utilized. This increases time and effort, and reduces efficiency.

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 SUPERDOME FLEX

The ideal solution for solving complex, data-intensive problems holistically is with HPE shared memory high performance computing and the breakthrough SMP (symmetric multiprocessing) system, HPE Superdome Flex. Providing unparalleled scale-up compute and shared memory resources with single-system simplicity, HPE Superdome Flex equips your HPC team to:

Accelerate time to discovery as jobs no longer compete and wait for "fat nodes", enabling more jobs to be completed in less time.

Improve accuracy using simulations with larger models and more data.

Increase productivity by freeing scientists and engineers from cluster management tasks.

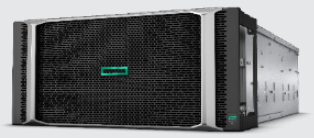

HPE Superdome Flex
4-socket, 5U
Modular Building Block

HPE Superdome Flex scales seamlessly as a single-system in 4-socket increments*

4 sockets (up to 6 TB†) | **8 sockets** (up to 12 TB†) | **16 sockets** (up to 24 TB†) | **20 sockets** (up to 30 TB†) | **32 sockets** (up to 48 TB†)

* 12-socket, 24-socket, and 28-socket configurations not pictured
† Amount of shared memory

- Solve complex, data-intensive problems **holistically** at **unparalleled** scale with **single-system** simplicity
- Complete **more jobs in less time**
- Free scientists and engineers from managing clusters to **accelerate time to discovery**

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. Superdome Flex is like a giant Linux® workstation with lightning speed.

This platform combines the best of HPE technologies, including from acquired high-performance compute innovator SGI, to deliver:

- **Unparalleled scale.** Through a unique modular architecture, the system scales seamlessly from 4 to 32 2nd generation Intel® Xeon® Scalable processors (Gold or Platinum) with up to 892 cores, and from 768 GB to 48 TB of shared memory, in a single system.
- **Optimum flexibility.** The 4-socket chassis building block design equips you with a system you will never outgrow, without having to overprovision. Memory capacity can be provided using DRAM, or a combination of DRAM and HPE Persistent Memory featuring Intel® Optane™ DC Persistent Memory technology. Add up to 16 NVIDIA® GPUs. Leverage 1/10/25/100 Gbe, 16/32 Gb FC, and IB EDR network connectivity, and internal SAS and NVMe SSDs, attached storage, and shared storage within a cluster using Multi-Rail LNet for Lustre.
- **Unbounded I/O.** Capitalizing on a high-bandwidth (13.3 GB/s) and low latency (<400ns) fabric providing massive data movement, utilize up to 16 LP/FH PCIe stand up cards per chassis and up to 128 cards per system.
- **Extreme availability.** Enjoy the highest service levels 24x7, enabled by unique HPE RAS features such as firmware-first that contains system errors at the firmware level and insulates the O/S; an error-analysis engine with auto-self healing; and advanced memory resiliency that helps prevent multi-bit memory errors from causing system failure. Couple Superdome Flex with HPE Serviceguard for Linux to provide maximum application availability. For a detailed list of all the HPE Superdome Flex RAS capabilities, refer to this technical white paper.
- **Simplified use experience.** Utilize HPE OneView Management, Insight Remote Support and OpenStack cluster management. Leverage industry-standard Redfish APIs to communicate management actions. HPE Superdome Flex provides a rich ecosystem of management tools and interfaces to ease the user experience. To learn more about HPE Superdome Flex server management, refer to this technical white paper.

REAL-WORLD RESULTS

Organizations from a wide range of industries are leveraging HPE Superdome Flex to tackle data-intensive HPC workloads holistically.

DZNE: Accelerating research
The large German research institute DZNE uses Big Data analytics to research and fight neurodegenerative diseases. With 1,000 employees in 80 working groups across 9 locations in Germany, DZNE was in the midst of a 30-year study tracking the health of 30,000 people when it reached the limits of traditional systems. DZNE overcame this bottleneck with an HPE in-memory HPC solution, achieving infinite scalability and a 100X anticipated increase in analytics speed.

"By storing a lot of data in memory, we have a much faster system that can accelerate our computational pipelines"

– Professor Joachim L. Schultze, Ph.D., Director, DZNE

NIG: Building a genomic supercomputer
The National Institute of Genetics is the center of life science and genomic medicine research in Japan, and provides researchers with opportunities for joint use of and research with supercomputers. NIG adopted Superdome Flex as its "5th generation supercomputer" to manage large-scale and complex environments in an integrated manner and respond quickly and efficiently to researchers' needs.

"A large shared memory space exceeding 10 TB is very effective for assembly processing that connects a large number of genome fragments to restore sequences"

– Dr Osamu Ogasawara, Project Associate Professor, NIG

COSMOS: Revealing the Big Bang
The UK's COSMOS advanced computing facility, co-founded by Stephen Hawking, runs complex simulations and real-time analyses of terascale datasets. The Superdome Flex computing platform is opening up new research horizons, helping COSMOS achieve its objectives of developing a seamless history of the Big Bang and understanding the gravitational waves of black holes.

"In-memory computing allows us to ingest all of this data and act on it immediately, trying out new ideas, new algorithms. It accelerates time to solution and equips us with a powerful tool to probe the big questions about the origin of our universe"

– Paul Shellard, Professor of Cosmology and Coordinator of COSMOS

Superdome Flex highlights

- Modular scale-up architecture with 4–32 sockets and 768 GB–48 TB of shared memory in single system
- Second generation Intel Xeon Scalable processors
- Start small and grow as needed with 4-socket chassis building blocks
- Unbounded I/O with up to 128 PCIe cards
- Up to 16 NVIDIA GPUs
- Extreme RAS for the highest service level demands
- Simplified user experience

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 with HPE Superdome Flex.

LEARN MORE AT

hpe.com/superdome

Make the right purchase decision.
Contact our presales specialists.

Chat **Email** **Call**

Get updates

Hewlett Packard Enterprise

a50000523ENW, August 2020, Rev. 1