

EVIDEN


BullSequana XH3500

Accelerate HPC simulations and AI factories with a converged supercomputer platform and its E2E solutions

One platform.

Any workload.

Unprecedented performance.



The rise of AI applications and increasingly complex simulations is pushing HPC to its limits. Customers now demand exascale-class performance, extreme density, technology flexibility, and energy efficiency.

BullSequana XH3500 meets these challenges with a modular, all-in-one hybrid supercomputer platform, combining large-scale AI and exascale simulation. It delivers **unmatched scalability, efficiency, flexibility, and resilience**, powered by Eviden's decades of R&D and supercomputing expertise.

Unmatched Scalability & Future Readiness

- **Granular scalability** from blade to rack with an **all-in-one** modular architecture for fine-grain expansion.
- **Future-ready** design with substantially increased power and cooling.

Best-in-Class Energy Efficiency

- **5th-generation 100% fanless Direct Liquid Cooling (DLC) operating at 40°C (104°F) facility water** to reach maximum efficiency and deliver lowest TCO.
- **Intelligent energy optimization software suites** with BullSequana ARGOS and Smart Performance Management Suite for real-time power capping, carbon tracking, workload scheduling, and GPU-aware tuning, maximizing performance per watt.

- **Dynamic resource optimization** through software that monitors job efficiency, allocates resources, and fine-tunes performance, maximizing utilization and reducing scale-out needs.

Industry-leading Density & Flexibility


- **High-density** architecture provides industry-leading compute density.
- **Full flexibility** with interchangeable CPU/GPU compute blades and interconnect switch blades, supporting any leading technology to avoid technology vendor lock-in.

Trusted, Sovereign, and Secured by Design


- **N+1 redundancy** for power and cooling, covering pumps and controllers, ensures high availability and continuity even under stress or failure conditions.
- Security is embedded at the hardware level, with features such as a **hardware root of trust, secure boot, and secure firmware update**, protecting the system against unauthorized modification and attacks.
- All BullSequana software includes a **Software Bill of Materials (SBOM)** to identify and remediate vulnerabilities throughout the cluster lifecycle, ensuring maximum security from development to production.
- Trusted **global support** tailored to operational needs, with options for 'Made in' regional delivery, ensuring sovereignty, reliability, and compliance worldwide.



Unmatched Scalability
& Future Readiness



Best-in-class
Energy Efficiency



Industry-leading
Density & Flexibility



Trusted, Sovereign,
and Secure by Design

Technical Specifications

BullSequana XH3500 Rack Specifications

Physical dimensions with doors on castors/loadspreader	Without Ultra Capacitor Module (UCM), H x W x D in mm (in): 2270 (89.4) x 900 (35.4) x 1457 (57.4) With UCM, H x W x D in mm (in): 2436 (95.9) x 900 (35.4) x 1457 (57.4)
Cabinet maximum weight at maximum rack configuration	Without UCM, in kg (lbs): 3211 (7079) With UCM in kg (lbs): 3449 (7604)
Cabinet maximum weight density on load spreader at maximum rack configuration	Without UCM, in kg/m ² (lbs/ft ²): 2744 (562) With UCM, in kg/m ² (lbs/ft ²): 2947 (603)
Power requirements	AC compute: up to 284kW (compute + 1 pump) DC compute: up to 264kW with N+1 PSU redundancy Optional UCM to provide power failure up to 300ms Supported Voltages: 400V (EU) 415V-480V (US) and 380V (Brazil-Korea) Rack power connections are standard top feed, bottom feed with accessories: • 4x 125A 230/400V (EU) 3P+N+E • 2x 16A 230/400V (EU) 3P+E (one active pump, one backup)
Cooling	100% fanless Direct Liquid Cooling, in-rack Coolant Distribution Unit. Rack inlet Water Temperature: 6 – 40 °C (43 – 104 °F) depending on site environmental configuration and blade technology. Compliance: ASHRAE A2/W40. Redundancy: Redundant parts include pumps, controllers, and primary regulation valves. Rack hydraulic connections: Standard bottom feed, top feed available with accessories. Connections: • 2x DN38 (Rack side) • 2x DN40 with insulated hoses (Customer side)
Blade slots	38 universal 1RU slots • House compute blades in different sizes: 1RU or 2RU • Support interconnect switch blades • Accommodate up to eight 1RU UCM blades

BullSequana XH3515B Accelerated Compute Blade Specifications

Form factor	1RU
Blade layout	Single node GB200-NVL4 (4X NVIDIA Blackwell B200 GPU, 2× NVIDIA Grace CPU)
GPU	NVIDIA Blackwell B200 186GB HBM3e per GPU
CPU	NVIDIA Grace 512GB LPDDR5x per CPU
Compute networking	4X Eviden BXI V3 or 4× InfiniBand NDR/XDR
I/O networking	Up to 1x InfiniBand NDR or 1x Ethernet
Node storage	Up to 9x E1.S NVMe SSD

BullSequana AI1242 Accelerated Compute Blade Specifications

Form factor	2RU
Blade layout	Single node (AMD Instinct™ MI355X GPU, 2 x AMD EPYC™ Turin CPU)
GPU	AMD Instinct™ MI355X 288GB HBM3e per GPU
CPU	AMD 5th Gen EPYC™ (Turin) 12x memory channels of DDR5 6400MT/s per CPU
Compute networking	8x Eviden BXI V3 or 8x InfiniBand NDR/XDR
I/O networking	Up to 2x BXI V3 or 2x InfiniBand NDR or 2x Ethernet
Node storage	Up to 4x E1.S NVMe SSD

Eviden's ecosystem: maximizing BullSequana XH3500 with an end-to-end approach



Integrated Solutions and Expertise

Eviden's HPC and AI ecosystem provides a fully integrated stack designed to accelerate innovation from real-world applications and use-cases to infrastructure.


At the top are the applications and industry use cases that deliver measurable impact across sectors such as energy, manufacturing, aerospace, healthcare, and telecommunications.

These are powered by Eviden's Center for Excellence in Performance Programming (CEPP), which fine-tunes code and applications to achieve optimal performance and efficiency.

The infrastructure layer brings together Eviden's high-performance BXI V3 interconnect, advanced software, and scalable storage, providing the throughput and reliability required by intensive HPC and AI workloads.

At the foundation lies the BullSequana XH3500 platform, delivering the compute power that fuels scientific simulation and AI.

Surrounding and supporting every layer, Eviden's Services and Support, combined with its Modular Data Center solutions, enable flexible deployment, rapid implementation, and faster time-to-market for every customer.



Accelerate HPC simulations and AI factories with a converged supercomputer platform and Eviden's ecosystem.

Connect with us



eviden.com

Eviden is a registered trademark © Copyright 2025, Eviden SAS – All rights reserved

