


BullSequana XH2000



The very nature of simulation is shifting from a traditional - monolithic scalar based computing - to a purpose-built workflow that may include any type of CPU, GPU, Vector units, TensorFlow units, clustered or shared memory computing, on premise up to the public cloud. It is of utmost importance that your supercomputer solution has all the flexibility to run seamlessly any type of workload, anywhere, now and tomorrow.

The new BullSequana XH2000 is a breakthrough evolution of the market leading BullSequana X1000 supercomputer, the answer to rising hybrid computing requirements.

With BullSequana XH2000, Atos brings to the market a no compromise, highly flexible, Exascale-ready and lowest TCO on the market, Hybrid Computing solution for today's and tomorrow's real-life problem solving.

A hybrid computing solution

Hybrid computing technologies

- Combine a broad variety of CPUs, accelerators and high-speed interconnect networks to run mixed HPC workloads or dedicate a full GPU and High-speed Ethernet-based system to Deep Learning.
- Run any type of workflow, on the same supercomputer, without compromising performance.

Hybrid computing environments

- Combined with Smart Management Center (SMC) software suite and Bull eXtreme Factory hub (XF hub), BullSequana XH2000 integrates seamlessly in secure hybrid cloud ecosystems, allowing optimal workload orchestration between on-premises, private and public cloud environments.

Efficient by design

Energy efficient

- 100% of the system's critical components are equipped with our cutting-edge Direct Liquid Cooling (DLC) technology, enabling a data center PUE as close as possible to 1.0.
- Smart Energy Management Suite optional software allows a fine-grained energy consumption monitoring, as well as dynamic power optimization, thus reducing the power envelope.

Application efficient

- Smart Data Management Suite optional add-ons enhances application performances by dynamically reducing IO related bottlenecks, without requiring any application change.
- Installed on BullSequana XH2000, the Atos optional Codex AI Suite empowers you to rapidly develop and deploy robust Machine Learning and Deep Learning applications. Atos AI experts will guide you through your projects.

Cost efficient

- Using XF hub, find the perfect balance between CAPEX - a BullSequana XH2000 on-premises system tailored to your everyday needs - and OPEX - on-demand competitive public or private cloud solutions.

Highly flexible

Modular and scalable

- Mix different types of current and future computing technologies (GPUs, CPUs) and interconnect networks (InfiniBand HDR, HDR100, BXL, High-speed Ethernet) within one system to create a supercomputer that matches perfectly your needs
- Select the desired network topology (Full Fat Tree or DragonFly+) and pruning ratio.
- BullSequana XH2000 scales from one-rack, up to exaflop systems.

Exascale-ready

A true evolution of BullSequana X1000

- BullSequana XH2000 keeps all high standard technological features that have made BullSequana X1000 a leader in the race for exascale, while embracing customers' ever evolving needs for technological convergence.
- BullSequana XH2000 introduces DragonFly+ topology support, allowing greater scalability.

A futureproof solution


We are committed to making BullSequana XH2000 a sustainable long-term investment

- We will keep on improving our solutions and integrating new cutting-edge technologies within BullSequana XH2000
- With the Open BullSequana program, Atos opens BullSequana XH2000 to 3rd party technology providers willing to integrate their value on this platform.

Benefit from Atos expertise

- With the Fast Start program, Atos experts guide you from day one to make sure that the solution is tailored to your needs and that you can start running workloads efficiently as quickly as possible.
- Atos Center for Excellence in Parallel Programming collaborates with you to get optimal application performance on BullSequana XH2000.

The BullSequana XH2000 flexible packaging




The BullSequana XH2000 cabinet

Dimensions /Weight	Mounting Capacity: 42U - Without UCM: 2020x750x1270 mm (79.5x29.5x50 inches) With UCM (max): 2400x750x1270 mm (94.5 x 29.5 x 50 inches) Without compute blade or UCM: 1,250 kg (2,756 lbs.) - Packing weight (without compute blade or UCM, with pallet): 1,230 kg (2,712 lbs.) - Max weight (fully equipped): 2,035 kg (4,487 lbs.) without UCM, 2135kg (4,707 lbs.) with UCM.
Power	The power section, located at the top of the cabinet, has a form factor of 11.5 U. It is composed of a power distribution unit (PDU), power supply unit (PSU) shelves, optional ultra capacitor module (UCM) and a busbar to distribute power to all the components within the cabinet. Power Distribution Unit It has a form factor of 2.5U. It protects and distributes power to the PSU shelves and HYCs. It also monitors the power circuit with a power management controller module. Power cables to be supplied to the PDU vary depending on the region of the world where the system is installed: • 3x 5G16 63A tri-phased 400V AC EU input power lines, or • 6x 4G16 63A tri-phased 208V AC US input power lines. Power Supply Unit Shelves Up to 6 x 15kW direct liquid cooled PSU shelves. Each shelf has a 1.5U form factor and contains up to 5x 3kW PSU blocks. The level and type of redundancy is selectable: • Level: redundancy at the PSU block or at the PSU shelf level • Type: N, N+1, N+2, and 2N. Ultra-Capacity Module UCM chassis is optional. For the mitigation of micro power outages up to 300ms at full load or 800ms at 45% load when 3-phase uninterruptible power supply equipment is not present upstream in the data center infrastructure. Busbar Busbar distributes power from the PDU to all the components within the cabinet. All components are hot-pluggable.
Cooling	BullSequana XH2000 has a fan-less design and a unique Direct Liquid Cooling (DLC) technology that uses warm water up to 40°C to cool all critical components within the cabinet (compute blades, interconnect and management switches, and PSUs). The cooling system is composed of hydraulic chassis (HYC), primary and secondary manifolds, and an expansion tank. Hydraulic Chassis: HYC contain the heat exchanger system that allows it to achieve 95% of heat transfer between the primary and secondary manifolds. Up to 3 HYC are available depending on the redundancy type desired (N or N+1). Third HYC is for N+1 redundancy only. Primary and secondary manifolds: The primary manifold system connects the customer water loop to the HYC primary water inlets. The secondary manifold system connects HYC outlets to each compute blade, each management and interconnect switch and each PSU in the cabinet. All critical components are direct liquid cooled and mount directly onto the secondary manifold via hydraulic non-spill quick disconnect couplings. Expansion tank: The expansion tank prevents excessive pressure within the hydraulic circuit.

Management Network	Up to 2 management network switches, located at the top rear of the cabinet: - Ethernet: up to 48 ports, 1 Gb/s or 10Gb/s, depending on customer's bandwidth requirements. These switches are linked to top switches located in service racks: - Ethernet: up to 48 ports, 1Gb/s or 10Gb/s depending on customer's bandwidth requirements.
Regulatory compliance / Certifications	Safety: EC, IEC, UL and CSA Electromagnetic Compatibility: EC, FCC, ICES-03 and VCCI Environment: RoHS II & WEEE directives, REACH regulation.
Warranty	Standard warranty: 1 year Extended Warranty: consult your local Sales representative.

Compute blades

BullSequana XH2000 cabinet can accommodate up to 32 direct liquid cooled blades (20 at the front and 12 at the rear of the cabinet). Each 1U blade contains a cold plate with active liquid flow which permits to cool off all critical components by direct contact.

	BullSequana X2415 Nvidia A100 blade 	BullSequana X2410 AMD blade 	BullSequana X1120 Intel blade 	BullSequana X1125 Nvidia V100 blade 
Design	1U blade with 1 accelerated compute node	1U blade with 3 compute nodes side-by-side	1U blade with 3 compute nodes side-by-side	1U blade with 1 accelerated compute node
Processors	2 AMD EPYC Rome/Milan processors 4 Nvidia A100 SXM4 GPUs	3x 2 AMD EPYC Rome/Milan Processors	3 x 2 Intel® Xeon® Processor Scalable Family	2 Intel Xeon Scalable Cascade Lake processors 4 Nvidia V100 SXM2 GPUs
Architecture	AMD SP3 Platform Nvidia HGX A100	AMD SP3 Platform	Intel C620 chipset	Intel C620 chipset
Memory	16 DDR4 memory slots (up to 128GB 3200MT/s DIMMs)	3x 16 DDR4 memory slots (up to 128GB 3200MT/s DIMMs)	3 x 12 DDR4 memory slots (up to 128GB 2933MT/s DIMMs) Optional Optane DCPMM memory available	12 DDR4 memory slots (up to 128GB 2933MT/s DIMMs)
I/O slots	2x 1 Infiniband HDR single NIC mezzanine board or 2x 1 Infiniband HDR dual NIC mezzanine board or 2x 1 BXL single NIC mezzanine board	3x 1 InfiniBand HDR/HDR100 single NIC mezzanine board or 3x 1 BXL single NIC mezzanine board	3x InfiniBand HDR/HDR100 single NIC mezzanine board or 3x BXL single NIC mezzanine board	2x 1 InfiniBand HDR100 single NIC mezzanine board or 2x 1 BXL single NIC mezzanine board
Storage	1 SATA M.2 drive or 1 NVMe M.2 drive	3x 1 SATA M.2 drive or 3x 1 NVMe M.2 drive	3x 1 SATA 2.5" SSD or 3x 1 NVMe U.2 SSD	1 SATA 2.5" SSD
Power supply	PSU shelves on top of XH2000 cabinet			
Cooling	Cooling by direct contact			
Physical specifications (HxWxD)	44.45 x 600 x 540 mm (1.75 x 23.6 x 21.3 inches)			
OS and software	Red Hat Enterprise Linux and Smart Management Center			
Warranty	Standard warranty: 1 year Extended warranty: consult us			
Regulatory compliance	Safety: EC, IEC, UL, CSA certifications; Electromagnetic Compatibility: EC, FCC, ICES-03, VCCI certifications; Environment: RoHS II & WEEE directives, REACH regulation			

Interconnect Network & Topologies

The Interconnect Network is composed of the Interconnect Network switches and a uniquely designed network connection mid-plane.

Interconnect Network switches	Up to 10 direct liquid cooled Interconnect Network switches. They are located at the top rear of the cabinet. 3 technologies are available: • InfiniBand HDR/HDR100: 200Gb/s (HDR) or 100 Gb/s (HDR100), 40 ports (HDR) or 80 ports (HDR100) • Bull eXascale Interconnect (BXI): 100 Gb/s, 48 ports • High-speed Ethernet: up to 100 Gb/s, up to 48 ports.
Network connection mid-plane	It is located at the center of the cabinet. It brings 3 major benefits: • Flexibility to customize routing of the compute blades to the interconnect network switches • Possibility to mix different interconnect network speeds and/or technologies • Selection of the optimized interconnect network topology (Full Fat Tree, DragonFly+).
Topologies	BullSequana XH2000 supports two topologies: • Full Fat Tree: a proven network architecture that provides very good worst-case blocking performance • DragonFly+: a cost-efficient topology that allows greater scalability than Full Fat Tree.

Software

Smart Management Center	SMC is the software cluster management stack for BullSequana XH2000 • SMC: up to 1200 computes nodes • SMC xScale: 1200 nodes and above, up to Exascale requirements with services orchestration.
Atos HPC added value software suites	• Smart Energy Management Suite: including energy optimizer and dynamic power optimizer features • Smart Data Management Suite: IO Instrumentation, FastIO Libraries, IO Pattern Analyzer and Smart Bust Buffer • Smart Performance Management Suite: Slurm workloads scheduler and utilities, Lightweight profiler, Open MPI and utilities, and Interconnect Bull Management Suite • eXtreme Computing Studio: modular HPC, AI and Quantum portal • eXtreme Remote Visualizer: 3D remote visualization software.
Hybrid computing add-ons	Bull eXtreme Factory hub: BullSequana XH2000 can be seamlessly integrated within a multi-cloud ecosystem thanks to extreme factory™ hub. This software layer supports HPC administrators in optimizing the workload dispatch between on-premise systems and private and public cloud solutions. Codex AI Suite : Machine Learning and Deep Learning are very powerful tools that can be challenging to implement. BullSequana XH2000's architecture can be optimized to support the Codex AI Suite, an environment that provides our customers with all necessary tools to easily define, develop and deploy robust AI applications.

For more information: hpc@atos.net

Atos, the Atos logo, Atos|Syntel, and Unify are registered trademarks of the Atos group. June 2020. © 2020 Atos. Confidential information owned by Atos, to be used by the recipient only. This document, or any part of it, may not be reproduced, copied, circulated and/or distributed nor quoted without prior written approval from Atos.

CT:200616-EM-FS-BullSequanaXH2000_DataSheet3