

# EVIDEN

**Leverage exceptional computing capabilities and strike a balance between cost and performance with the BullSequana AI mid-range server family.**

# BullSequana AI 600H, AI 600, AI 200 series

## Truly designed for any AI workload

The BullSequana mid-range AI server family enables you to handle various AI workloads and scale as your needs grow. These servers assist developers, data scientists, and researchers make significant advancements in innovation and computation. They provide exceptional computing capabilities and strike a balance between cost and performance, ensuring that customers receive the best return on their investment.

The BullSequana AI 200 and AI 600 server families feature the latest technologies, thanks to Eviden's close collaboration with key technology providers.

BullSequana mid-range AI server family and its ranges:

- BullSequana AI 600H will be available in liquid cooling configurations for energy efficiency and density in data centers.
  - Ranges: BullSequana AI 641-H01, BullSequana AI641-H02, BullSequana AI642-H01.
- BullSequana AI 600 is a medium-sized enterprise family requiring 8 top GPU servers.
  - Ranges: BullSequana AI 621-A01, BullSequana AI 622-A01, BullSequana AI 641-A01, BullSequana AI 641-A02 and BullSequana AI 642-A01.
- BullSequana AI 200 is a generic and flexible server family for enterprises that require up to 4 GPUs in PCIe format.
  - Ranges: BullSequana AI 221-A01 and BullSequana AI 222-A01.

## Experience unparalleled benefits and the most advanced technologies available in the market

- Latest market-leading high-performance computing technologies (Intel®, NVIDIA® and AMD®) both in PCIe format and 8 GPUs board format with high-speed interconnection between the GPUs.
- Largest DIMM and NVDIMM capacity at high-frequency rates.
- High-end interconnect networking technologies (Nvidia InfiniBand, High-Speed Ethernet).
- Wide storage technologies (NVME, SATA).

All our servers are supported by Eviden HPC-AI Software Suites, a set of software solutions that fulfill the demands of even the most challenging HPC configurations and workloads while maintaining high-security standards.

## Take advantage of Eviden's extensive experience in AI solutions

With the BullSequana AI mid-server family, Eviden offers outstanding service. By choosing the BullSequana AI, you receive guidance and expertise from highly qualified AI professionals.

- With the Fast Start program, our team of experts will personally guide you from day one to ensure that the solution is tailored to your specific needs and that you can start running workloads efficiently as quickly as possible.
- Our highly skilled engineers will install and evaluate your AI supercomputer on your premises, addressing any questions you may have.
- You'll receive high-quality maintenance and support throughout the supercomputer's lifecycle.
- The Eviden Center for Excellence in Performance Programming will collaborate with you to achieve optimal application performance.
- With the Eviden Managed Services option, we can efficiently manage and operate your state-of-the-art AI supercomputers.

## Embraced an optimized Total Cost of Ownership (TCO)

The BullSequana AI mid-range server family has been meticulously engineered to deliver cost-efficiency alongside uncompromised performance. Each server within is finely tuned for optimized density and an exceptional performance-to-cost ratio. BullSequana AI mid-range servers are air-cooled and liquid-cooled, fitting in standard racks and enabling quick deployment while saving space and cost.

BullSequana AI's mid-range family is highly energy efficient. It features 80+ titanium-certified Power Supplies, PWM cooling fans, and advanced power monitoring and management via Eviden HPC Software Suites.

Numerous customers have trusted Eviden BullSequana AI mid-range servers in government, education, and industrial facilities.

### **BullSequana AI 640 and BullSequana AI 640H**

BullSequana AI 640 servers are perfectly sized to deliver the performance required by power-hungry and complex AI workloads. Their outstanding density allows them to minimize their footprint, significantly reduce real estate or rental costs, and improve energy efficiency. This server is Nvidia AI-certified and is especially suited for large LLM models.

BullSequana AI 640H allows best-in-class TCO and density in the best value-for-money solution.

### **BullSequana AI 620**

BullSequana AI 620 provides comprehensive GPU-based solutions, integrating the latest hardware technologies, software environment, expertise, and services for seamless implementation of your solution.


NVIDIA® and AMD GPU are highly parallel processors designed to boost computing. Integrating up to eight GPUs within 4U places this system at the forefront of today's accelerated computing solutions in compute density, performance, and cost optimization.

Eviden pioneered GPU accelerators in 2008 in its supercomputing range. They now power energy-efficient data centers in government labs, universities, enterprises, and small and medium businesses worldwide.

### **BullSequana AI 200**

BullSequana AI 200 is a highly flexible GPU node and affordable entry solution that combines the qualities expected from efficient nodes in the most demanding AI, data visualization, and rendering workloads.




The server is optimized to support multiple coprocessors, GPUs, and the latest graphics cards. It is compatible with a wide range of NVIDIA® and AMD GPU cards for computing and visualization purposes.



# BullSequana AI 600H, AI 600 and AI 200 series


BullSequana AI 600 series is a server range supporting the latest NVIDIA and AMD accelerators. NVIDIA and AMD accelerators are based on the latest technology and offer the highest density by scaling from 1 to 8 accelerators. The BullSequana AI, 600 family of servers, is available in dual sockets, with memory support from 8 to 24 DDR5 channels per socket from 4800MT/s to 6000 MT/s. Each node allows the fastest interconnect and GPU direct possibility. The BullSequana AI600H series is the DLC version for better energy efficiency and density in data centers. Both PCIe format and board format are available. Nvidia SXM format with NVlink and AMD OAM format with Infinity fabric allow the best AI workload and LLM performances.

## Technical specifications

	BullSequana AI 641-A01	BullSequana AI 641-H01
		
Cooling	Air	Liquid
Form Factor	5U 1N 2S	
Dimensions HxWxD, mm	222.25 x 447 x 945	
GPU Accelerator and Graphics	NVIDIA HGX H100	
Processors	Dual Socket 4th and 5th Gen Intel® Xeon® processors	
Memory	32x DIMM slots DDR5 4800/5600MHz RDIMM/LRDIMM	
Expansion slots per server	8x LP PCIe 5.0x16 4x FHHL PCIe 5.0x16	
Network per server	2x 10GbE ports 1x Management port	
Storage	Front	8 x 2.5" Gen5 NVMe/SATA/SAS
Power Supply	Redundant 4+2 3000W 80+ Titanium	
OS & Cluster Software	Red Hat Enterprise Linux & Eviden HPC Software Suites Electromagnetic Compatibility (FCC, CE, ICES-03 certification) Environment (RoHS II & WEEE directives)	
Regulatory Compliance	Safety (CE certification)	

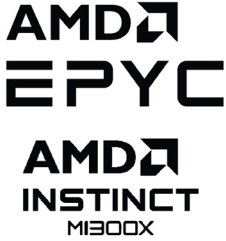

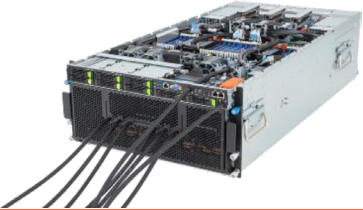
# BullSequana AI 600 / AI 600H series

## Technical specifications

	BullSequana AI 641-A02	BullSequana AI 641-H02
Cooling	Air	Liquid
Form Factor	5U 1N 2S	
Dimensions HxWxD, mm	222.25 x 447 x 945	
GPU Accelerator and Graphics	NVIDIA HGX H200	
Processors	Dual Socket	
	4th and 5th Gen Intel® Xeon® processors	
Memory	32x DIMM slots	
	DDR5 4800/5600MHz RDIMM/LRDIMM	
Expansion slots per server	8x LP PCIe 5.0x16 4x FHHL PCIe 5.0x16	
Network per server	2x 10GbE ports 1x Management port	
Storage	Front	8 x 2.5" Gen5 NVMe/SATA/SAS
Power Supply	Redundant 4+2 3000W	
	80+ Titanium	
OS & Cluster Software	Red Hat Enterprise Linux & Eviden HPC Software Suites- Electromagnetic Compatibility (FCC, CE, ICES-03 certification) - Environment (RoHS II & WEEE directives)	
Regulatory Compliance	Safety (CE certification)	
Warranty & Services	1-year standard warranty, optional extension	



# BullSequana AI 600 / AI 600H series

## Technical specifications

	BullSequana AI 642-A01	BullSequana AI 642-H01
		
Cooling	Air	Liquid
Form Factor	5U1N2S	
Dimensions HxWxD, mm	222.25 x 447 x 945	
GPU Accelerator and Graphics	AMD MI300X	
Processors	Dual Socket	
	4th and 5th Gen AMD EPYC™ processors	
Memory	24x DIMM slots	
	DDR5 4800/6000MHz RDIMM/LRDIMM	
Expansion slots per server	8x LP PCIe 5.0x16 4x FHHL PCIe 5.0x16	
Network per server	2x 10GbE ports 1x Management port	
Storage	8 x 2.5" Gen5 NVMe/SATA/SAS	
Power Supply	Redundant 4+2 3000W	
	80+ Titanium	
OS & Cluster Software	Red Hat Enterprise Linux & Eviden HPC Software Suites Electromagnetic Compatibility (FCC, CE, ICES-03 certification) Environment (RoHS II & WEEE directives)	
Regulatory Compliance	Safety (CE certification)	
Warranty & Services	1-year standard warranty, optional extension	



# BullSequana AI 600 series

## Technical specifications

 <p><b>AI Enterprise   AI certified</b></p>	<b>BullSequana AI 621-A01</b>	
		
<b>Cooling</b>	Air	
<b>Form Factor</b>	4U1N2S	
<b>Dimensions HxWxD, mm</b>	176 x 448 x 880 mm	
<b>GPU Accelerator and Graphics</b>	NVIDIA H100, H100 NVL, L40, L40s, L4 AMD Radeon MI210	
<b>Processors</b>	Dual Socket 4th and 5th Gen Intel® Xeon® processors	
<b>Memory</b>	32x DIMM slots DDR5 4800/5600MHz RDIMM/LRDIMM	
<b>Expansion slots per server</b>	8x FHFL PCIe 5.0x16 8x LP PCIe 5.0x16 2x LP PCIe 5.0x16	GPU direct compatible
<b>Network per server</b>	2x 1GbE ports 1x Management port	
<b>Storage</b>	Front	Up to 12x 2.5" SATA/U.2 disks
<b>Power Supply</b>	Redundant 3+1 3000W 80+ Titanium	
<b>OS &amp; Cluster Software</b>	Red Hat Enterprise Linux & Eviden HPC Software Suites	
<b>Regulatory Compliance</b>	Safety (CE certification) Electromagnetic Compatibility (FCC, CE, ICES-03 certification) Environment (RoHS II & WEEE directives)	
<b>Warranty &amp; Services</b>	1-year standard warranty, optional extension	

# BullSequana AI 600 series



## Technical specifications

	<b>BullSequana AI 622-A01</b>	
		
<b>Cooling</b>	Air	
<b>Form Factor</b>	4U1N2S	
<b>Dimensions HxWxD, mm</b>	176 x 448 x 880 mm	
<b>GPU Accelerator and Graphics</b>	NVIDIA H100, H100 NVL, L40, L40s, L4 AMD Radeon MI210	
<b>Processors</b>	Dual Socket	
	4th and 5th Gen AMD EPYC™ processors	
<b>Memory</b>	48x DIMM slots	
	DDR5 4800/6000MHz RDIMM/LRDIMM	
<b>Expansion slots per server</b>	8x FHFL PCIe 5.0x16 8x LP PCIe 5.0x16 2x LP PCIe 5.0x16	
<b>Network per server</b>	2x 1GbE ports 1x Management port	
<b>Storage</b>	Front	Up to 12x 2.5" SATA/U.2 disks
	Internal	2x M.2 2280/22110
<b>Power Supply</b>	Redundant 3+1 3000W	
	80+ Titanium	
<b>OS &amp; Cluster Software</b>	Red Hat Enterprise Linux & Eviden HPC Software Suites	
<b>Regulatory Compliance</b>	Safety (CE certification) Electromagnetic Compatibility (FCC, CE, ICES-03 certification) Environment (RoHS II & WEEE directives)	
<b>Warranty &amp; Services</b>	1-year standard warranty, optional extension	





# BullSequana AI 200 series

## Technical specifications

	<b>BullSequana AI 221-A01</b>	
		
<b>Cooling</b>	Air	
<b>Form Factor</b>	2U1N2S	
<b>Dimensions HxWxD, mm</b>	87.5 x 438 x 815 mm	
<b>GPU Accelerator and Graphics</b>	NVIDIA H100, H100 NVL, L40, L40s, L4 AMD Radeon MI2100	
<b>Processors</b>	Dual Socket	
	4th and 5th Gen Intel® Xeon® processors	
<b>Memory</b>	32x DIMM slots	
	DDR5 4800/5600MHz RDIMM/LRDIMM	
<b>Expansion slots per server</b>	4x FHFL PCIe 5.0x16 for GPU 1x FHHL PCIe 5.0x16 1x OCP 3.0 PCIe 5.0x16	
<b>Network per server</b>	2x 1GbE ports 1x Management port	
<b>Storage</b>	Front	12x hot-swap 3.5"/2.5" SATA/NVMe
	<b>Power Supply</b>	Redundant 1+1 2700W
80+ Titanium		
<b>OS &amp; Cluster Software</b>	Red Hat Enterprise Linux & Eviden HPC Software Suites	
<b>Regulatory Compliance</b>	Safety (CE certification) Electromagnetic Compatibility (FCC, CE, ICES-03 certification) Environment (RoHS II & WEEE directives)	
<b>Warranty &amp; Services</b>	1-year standard warranty, optional extension	

# BullSequana AI 200 series

## Technical specifications

	<b>BullSequana AI 222-A01</b>	
		
<b>Cooling</b>	Air	
<b>Form Factor</b>	2U1N2S	
<b>Dimensions HxWxD, mm</b>	87.5 x 438 x 815 mm	
<b>GPU Accelerator and Graphics</b>	NVIDIA H100, L40, L4 AMD Radeon MI210	
<b>Processors</b>	Dual Socket	
	4th and 5th Gen AMD EPYC™ processors	
<b>Memory</b>	24x DIMM slots	
	DDR5 4800/6000MHz RDIMM/LRDIMM	
<b>Expansion slots per server</b>	2x FHFL PCIe 5.0x16 2x FHFL PCIe 5.0x16 1x FHHL PCIe 5.0x16 1x OCP 3.0 PCIe 5.0x16	
<b>Network per server</b>	2x 1GbE ports 1x Management port	
<b>Storage</b>	Front	4x hot-swap 3.5"/2.5" SATA 8x hot-swap 3.5"/2.5" SATA/U.2
<b>Power Supply</b>	Redundant 1+1 2700W	
	80+ Titanium	
<b>OS &amp; Cluster Software</b>	Red Hat Enterprise Linux & Eviden HPC Software Suites	
<b>Regulatory Compliance</b>	Safety (CE certification) Electromagnetic Compatibility (FCC, CE, ICES-03 certification) Environment (RoHS II & WEEE directives)	
<b>Warranty &amp; Services</b>	1-year standard warranty, optional extension	

Connect with us



**eviden.com**

Eviden is a registered trademark © Copyright 2023, Eviden SAS – All rights reserved.

ECT-240618-AR-BR-BullSequana AI 600-200-En-V1