

BullSequana S series

Technical specifications

To tackle enterprise IT challenges and enable businesses to take full advantage of Artificial Intelligence (AI), Atos brings to the market a new generation of x86 servers, BullSequana S, optimized for machine-learning, business critical computing applications and in-memory environments. BullSequana S reaches the highest level of quality of service, performance, availability and scalability to meet IT departments' existing and emerging demands.

Accelerating digital transformation

Powering Enterprise Artificial Intelligence

In order to utilize the extensive capabilities of AI, businesses require an infrastructure with extreme performance. BullSequana S meets the challenge with its unique combination of the most advanced Intel® Xeon® Scalable processors (CPUs) and Graphics Processing Units (GPUs). This innovative architecture, designed by Atos' R&D teams, enables to mix GPU, storage and compute modules within a single server for ready availability of all workloads.

Boosting data analytics & data lake

BullSequana S pre-integrated platform speeds up data lake environments deployments. Its ability to scale as needs arise and adjust with a vast array of internal disks, enables to finely tune the platform to the business requirements and simplify the lifecycle of new application.

Accelerating in-memory applications

The key element to go real-time is to have all structured data staged in-memory which requires to encompass specific



features. BullSequana S superior scalability, availability and serviceability makes it the ideal scale-up platform for very large enterprise applications and in-memory computing.

Optimizing IT modernization through virtualization & cloud

BullSequana S is the most agile, scalable and open platform to grow digital business. With its dynamic reconfiguration capabilities, it combines exceptional performance with unparalleled agility and generates efficiencies at every level. It is the go-to server for the private part of an hybrid cloud.

BullSequana S reconciles the extremes

Exceptional scalability - 2 to 32 CPUs / 48TB RAM / up to 32 GPUs.

To preserve investments and power the most demanding environments.

Modularity and flexibility

To build or reconfigure, fitting exactly your business needs.

Operations and TCO benefits for all landscapes

Versus classical environments, up to 20% cost reduction for large clusters virtualization and up to 30% price/performance gain for small & medium SAP HANA landscapes.



Atos

Technical specifications

	S200	S400	S800
DESIGN			
Form Factor	19" 2U	19" 4U	19" 8U
PROCESSORS			
Numbers	2 max 56 cores and 112 threads	2 - 4 max 112 cores and 224 threads	2 - 8 max 224 cores and 448 threads
Type	Intel® Xeon® Scalable processor - 8100, 6100, 5100, 4100 series - 4, 8, 12, 14, 18, 22, 24, 26 or 28 cores		
L3 shared cache	77MB	154MB	308MB
ARCHITECTURE			
Chipset	Intel® C627 Chipset		
Ultra-Path Interconnect	Intel® UPI: 2-3 links per socket - up to 10.4 GT/s		
Scalability	2 processors	2 to 4 processors	2 to 8 processors
Hardware Partitioning	No	Yes	Yes
MEMORY			
Memory Slots	24	48	96
Min / max	64 GB - up to 3 TB	128 GB - up to 6 TB	256 GB - up to 12 TB
Type	DDR4 RDIMM, LR-DIMM (64 & 128 GB only)		
EMBEDDED I/O PORTS			
Network Interface Controller (NIC)	4x 10Gb/S Copper Ethernet ports	8x 10Gb/S Copper Ethernet ports	16x10Gb/S Copper Eth. ports
Management ports	100 Mb/s	100 Mb/s with Private Ethernet management switch for 4S & 8S	
USB ports	4x USB 3.0 (3x front + 1x internal) + 1 x USB 2.0	4x USB 3.0 (3x front + 1x internal) + 1 x USB 2.0	4x USB 3.0 (3x front + 1x internal) + 1 x USB 2.0
I/O			
I/O slots	Up to 7 PCIe Gen3 slots: 5x PCIe blades x8 (4+1 dedicated for SAS controller per 2 CPUs) + 2x internal PCIe Gen3 x16 (including GPUs post RTS)	Up to 14 PCIe Gen3 slots: 10x PCIe blades x8 (4+1 dedicated for SAS controller per 2 CPUs) + 4x internal PCIe Gen3 x16 (including GPUs post RTS)	Up to 28 PCIe Gen3 slots: 20x PCIe blades x8 (4+1 dedicated for SAS controller per 2 CPUs) + 8x internal PCIe Gen3 x16 (including GPUs post RTS)
I/O availability	Hot swap PCIe blades depending on OS / hypervisor		
NIC PCIe blade	1GbE: 2 or 4 ports per PCIe blade ; 10GbE: 2 ports per PCIe blade		
HBA PCIe blade	8Gb/s: 2 ports per PCIe blade; 16Gb/s: 2 or 4 ports per PCIe blade		
SAS PCIe blade	12Gb/s: 2 ext. ports per PCIe blade		
STORAGE			
Internal disks Chassis	Front hot-swappable disks		
Compute Unit	Up to 8x 2.5" SSD/HDDs.	Up to 16x 2.5" SSD/HDDs.	Up to 32x 2.5" SSD/HDDs.
Storage Unit*	up to 12x 2.5" SSDs/HDDs or up to 4x 3.5" HDDs or up to 4x 2.5" NVMe U.2 drives	up to 24x 2.5" SSDs/HDDs or up to 8x 3.5" HDDs or up to 8x 2.5" NVMe U.2 drives	up to 48x 2.5" SSDs/HDDs or up to 16x 3.5" HDDs or up to 16x 2.5" NVMe U.2 drives
Storage Controller	RAID controllers: 9361-8i (SAS 12Gb/s, internal disks) , 9361-16i (Storage Unit) SAS 12Gb/s HBA: SAS9300-8i (internal disks), SAS9305-16i (Storage Unit)		
Micro SD	Dual MicroSD (RAID) on internal USB port		
SAN	Dell EMC, NetApp, ...		

* The 2-socket module hosts the Compute Unit and optionally either the Storage Unit or the GPU Unit in the top U.

	S200	S400	S800
VIDEO			
Video controller	1		
Memory	8 MB		
GPUs: GPU Unit*	Up to 2 GPUs	Up to 4 GPUs	Up to 8 GPUs
SECURITY			
Security features	TPM 2.0, Secure boot, 2-level password		
POWER SUPPLY			
Hot swap PSU	1 + 1 per module		
Power supply number	2, redundant	Up to 4, redundant	Up to 8, redundant
Power type	Label 80+ Titanium & Platinum, 96% efficient		
Power consumption	2000 watts		
Auto-sensing	220V 60/50Hz		
COOLING			
Fan specifications	Up to 14 hot swap, N+1 redundant	Up to 28 hot swap, N+1 redundant	Up to 56 hot swap, N+1 redundant
PHYSICAL SPECIFICATIONS			
Dimensions (HxLxW) (max)	89 (2U) x 446 mm (19") x 850 mm	175 (4U) x 446 mm (19") x 850 mm	352 (8U) x 446 mm (19") x 850 mm
	Recommended to use BullSequana racks		
Weight	43 kg maximum	Up to 81 kg	Up to 160 kg
Operating constraints	10°C at 35°C, gradient 20°C/h, 20% at 60%, gradient 5%/h		
OS & SOFTWARE			
Operating System	VMware® vSphere (ESXi™), Red Hat® Enterprise Linux®, Suse® Linux Enterprise Server, Microsoft® Windows Server		
SYSTEM MANAGEMENT			
BMC	IPMI 2.0		
Remote management	Standard via on-board iBMC (connection through the management port)		
Management software	BMC (Server Hardware Console), IPMI 2.0, RedFish, Bull iCare, Bull Platform Manager		
AVAILABILITY & RAS FEATURES			
RAS features	Advanced Error Detection and Correction (AEDC), Viral Mode of error containment, PCIe "Stop and Scream", Virtual (soft) Partitioning, PCI Express ECRC, PCIe Corrupt Data Containment (Data Poisoning), PCIe Link CRC Error Check and Retry, PCIe Link Retraining and Recovery, PCI Express Live Error Recovery, DDR4 Wr Data CRC check/retry, DDR4 Command/Address Parity Check and Retry, Intel® UPI Link Level Retry, Intel® UPI Protocol Protection via 16 bit Rolling CRC, Intel® UPI Dynamic Link width reduction, Core disable for Fault Resilient Boot, Power up, Post Package Repair, Failed DIMM Isolation, PCIe Card Hot Plug (Add/Remove/Swap), PIROM for System Information Storage		
Serviceability	Hot swap devices: PCIe blades (depending on OS), disks, power supplies, fans, front access disks, compute box design		
Redundancy	Power supplies, fans, disks with RAID		
WARRANTY & SERVICES			
Standard warranty	3 Years CRU		
Warranty extension	Bull Global Care		
Other services	IT infrastructure Advisory and energy audits Service Assurance HA, capacity and performance management Installation and integration services		
REGULATOR & SAFETY			
Conformity	Safety (CE, IEC, UL, CSA + APAC certifications), Electromagnetic Compatibility (EC, FCC, ICES-03, VCCI certifications), Environment (RoHS II & WEEE directives, REACH regulation)		

* The 2-socket module hosts the Compute Unit and optionally either the Storage Unit or the GPU Unit in the top U.
Please contact your local Atos representative for technical information on large configurations with BullSequana S1600 and S3200.

A powerful and evolutive range ...

Based on a very flexible architecture, the BullSequana S server range consists of 5 complementary models, based on the Intel® Xeon® Scalable processor. Each model can be smoothly upgraded to another one, preserving investments and applicative environments. This easy upgrade path is made possible thanks to a very modular design:

- A server scales from 2 to 32 processors, up to 32 GPUs with a maximum memory capacity of 48 TB RAM and 64 TB NV-RAM. All those components

within a server are hosted within 1 to 4 Compute boxes;



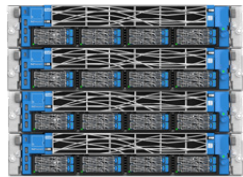


- The Compute box is the base element of a server, with 3 different form factors (2U/4U/8U) hosting 1 module per 2U;
- The interconnections are performed within a single Compute box by a Connecting box. Above 8 CPUs, a Ubox is added, hosting the new generation of Atos eXternal Node Controller (XNC).

The UBox enables to interconnect up to 4 Compute boxes, to form an SMP system

(Symmetric Multi-Processor) with up to thirty-two processor sockets in a CC-NUMA architecture.

The UBox is a 5U chassis. This VLSI-type (Very Large Scale Integration) integrated circuit is derived from technologies developed for mainframe servers and tuned for High Performance Computing.

Up to 8 CPUs, BullSequana S scales in glueless mode; to reach 16 CPUs, a UBox is added to interconnect 2 Compute boxes (2*8 CPUs) and 2 UBoxes are necessary to reach 32 CPUs.

BullSequana S200	BullSequana S400	BullSequana S800	BullSequana S1600	BullSequana S3200
				
2 CPUs, up to 3 TB RAM, up to 2 GPUs	up to 4 CPUs, up to 6 TB RAM, up to 4 GPUs	up to 8 CPUs, up to 12 TB RAM, up to 8 GPUs	up to 16 CPUs, up to 24 TB RAM, up to 16 GPUs	up to 32 CPUs, up to 48 TB RAM, up to 32 GPUs

... based on modules for optimal flexibility

The module is the building block of BullSequana S servers. Each module can be extracted from the Compute box for easy maintenance. The module comprises a compute unit plus a Storage unit or a GPU unit, in option.

Compute unit

With 2 CPUs, up to 3TB RAM and NV-RAM capabilities (up to 7.5TB), it includes up to 8 disks (2"5) and hot plug/swap PCIe blade for easy maintainability. Its 2U form-factor makes it a dense solution and there is an open space to add some optional tray for optimized flexibility: a Storage or GPU unit.



Compute unit with CPU/RAM/IO disks

GPU unit, for artificial intelligence

This option will be used mainly to introduce up to 32 GPUs in a single server in a very flexible way, 2 GPUs per module. Real-time algorithms and machine learning will use this huge processing power to run.



Compute unit + 2 GPUs (GPU unit)

Storage unit, for data-extensive needs

This unit can hold up to: 12 SAS/SSD 2"5 disks; 4 NL-SAS high capacity 3"5 disks; 4 NV-Mem for high I/O throughput. Thanks to this additional Storage unit, the capacity goes up to 20 disks in a 2U form-factor, and more than 2 raw PB in a 32 CPU server. This will be used in various use-cases from data lake to virtualization.



Compute unit + extended storage (storage unit)

Find out more about us
atos.net/BullSequanaS

Intel, the Intel logo, Xeon, and Xeon Inside are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.

© Atos November 2017 - All trademarks are the property of their respective owners. Atos, the Atos logo, Atos Codex, Atos Consulting, Atos Worldgrid, Bull, Canopy, equensWorldline, Unify, Worldline and Zero Email are registered trademarks of the Atos group. Atos reserves the right to modify this document at any time without notice. Some offerings or parts of offerings described in this document may not be available locally. Please contact your local Atos office for information regarding the offerings available in your country. This document does not represent a contractual commitment.