

Powering intelligence at the edge

The exponential growth of intelligent sensors and devices is generating an unprecedented amount of data. This is reshaping IT architectures, as increasingly powerful processing and machine learning inference capabilities are required at the Edge of the networks to enable next generation, transformative AI and IoT applications. BullSequana Edge has been designed to meet these challenges, delivering powerful AI inference and streaming analytics capabilities while ensuring that all data remains safe and secure.

Designed to operate outside of the datacenter

In order to reduce latency and to optimize network bandwidth, Edge servers often need to be placed in close proximity of IoT devices such as cameras, sensors or production robots. BullSequana Edge is capable to function in weather protected locations which are only partially temperature controlled and can therefore operate in a variety of locations such as airports, shops or factory floors. Moreover, BullSequana Edge offers very flexible deployment choices such as desktop, wall or rack mount options.

Hyperconverged infrastructure for accelerated Edge Data Analytics

Industry 4.0 applications require Edge Analytics with the lowest possible latency as well as rock solid data persistence. BullSequana Edge has therefore been designed to enable open source based hyperconverged infrastructure solutions which enable flexible resource sharing between nodes, centralized management and security hardening.

Streaming analytics solutions such as Spark and Kafka can flexibly be deployed on this stack. Through its powerful GPUs BullSequana Edge also supports accelerated machine learning algorithms enabled by RAPIDS and similar frameworks.



High end AI and machine learning performance

BullSequana Edge has been designed to provide leading AI acceleration capabilities for resources hungry Video streaming analytics. The server can host up to two powerful Nvidia Tesla T4 GPUs or optional FPGAs. This enables the inference of complex AI models right at the edge with lowest possible latency. Together with its powerful 16 core Intel® Xeon® processor, BullSequana Edge provides an outstanding compute power-pack for the implementation of up to 512 GB RAM and 28 TB storage.

WiFi, 4G, LTE and Lora

Thanks to its Wifi, 4G, LTE and Lora radio capabilities, BullSequana Edge also ensures IoT connectivity. The radios capabilities make the server independent from traditional network connectivity and allow deployment in locations which are not equipped for isolated.

Security and Data protection

When located outside of a datacenter, the risk of physical intrusion increases substantially. BullSequana Edge has therefore been equipped with an Intrusion Sensor which disables the machine in case of physical attacks. A secure boot process can be put in place including signed firmware, bootloader and the OS, protected by a FIPS 140-2 certified TPM and encrypted disks.

BullSequana Edge Key benefits

1. Designed for AI
2. Secured IoT & Cloud capabilities
3. End to end service from Atos.

BullSequana Edge Hardware Specifications

The BullSequana Edge is composed with the following elements:

Processor

CPU	1 socket 16 cores / 32 threads
Type	Intel® Xeon® D 2187NT or 2183IT
Vector Extension	AVX-512 / 1 FMA
Frequency	Processor base frequency 2.00 or 2.20 referring to CPU type GHz Max turbo frequency 3.00 GHz

GPU	Up to 2 Nvidia Tesla 4 Nvidia NGC Certified
-----	--

Architecture

Chipset	System On Chip (SOC)
---------	----------------------

Memory

Min / max	From 32 to 512 GB
Type	RDIMM & LR-DIMM 2666MT/s

Storage

Storage capacity (up to 2 discs)	SSD 480GB, 960GB or 1.92TB HDD 8TB or 14TB
----------------------------------	---

Embedded I/O ports

Network Interface Controller (NIC)	2 x 10Gb/s SFP+ (optional) 2 x 1Gb/s RJ45
System Management	<ul style="list-style-type: none"> 1Gb/s RJ45 BMC WIFI adapter (optional) OpenBMC REST API and RedFish support IPMI (optional - deactivated by default)
USB ports	2 x USB 3.0

Security

Security features	Intrusion detection switch TPM 2.0 (FIPS 140-2, EL4+) Secure boot (optional) Disk encryption (optional) Intel QAT* (IPSEC acceleration) Atos signed firmware - HSM Evidian
-------------------	--

I/O

I/O	PCI slots: <ul style="list-style-type: none"> 2 x PCIe Gen3 16 2x mPCI 2 SATA disks (SSD or HDD)
Accelerators (optional)	GPU <ul style="list-style-type: none"> Up to 2x NVIDIA T4 or FPGA Up to 2x FPGA 75W Up to 1x FPGA 150W

Power

Power	Power supply type: 500W Maximum power consumption: 380W
-------	---

I/O continued

Connectivity	<ul style="list-style-type: none"> mPCI adapters / M.2 converter WIFI dual band (2.4GHz - 5 GHz) LoRaWAN 3G/4G/LTE
Mounting options	<ul style="list-style-type: none"> Stand alone 19" 2U Rack mount kit (optional) DIN rail (optional) VESA Wall mount kits (optional)
Dimensions	430 x 290 x 86 mm (L/W/H)

Operational conditions

Standard operating temperature	5°C to 45°C with 2x T4 GPUs (41°F to 113°F)
Extended operating temperature	0°C to 48°C without GPU (32°F to 118°F)
Relative humidity	5% to 95% RH
Absolute humidity	During operation: 1 to 29 g/m3 Storage / transport: 1 to 29 g/m3
Air pressure in operation	70 to 106 kPa (conditions in mines are not considered)
Movement of surrounding air in operation	5 m/s
Radiations in operation	Solar: 700 W/m2 Heat: 600 W/m2
Operating Shock	30g half sine, 11 ms duration, 18 events on all 3 axes, IEC 60068-2-27
Non-operating/Transit Shock	50g half sine, 11 ms duration, 18 events on all 3 axes, IEC 60068-2-27
Operating Vibration	Frequency 5-500Hz, 3Grms, IEC 60068-2-64
Non-Operating/transit vibration	Frequency 5-500Hz, 5Grms, IEC 60068-2-64

Ecosystem and certifications

Software ecosystem certifications	OS <ul style="list-style-type: none"> certified: RedHat, VMWare ESXi, Microsoft 2019 Server tested: Ubuntu, Debian, CentOS Cloud connectivity certifications: MS Azure IOT, AWS Greengrass, NVIDIA EGX, Siemens Mindsphere Open Connect
Certification and Compliance for FINAL System	UL/CSA CE(RED, RoHS) FCC/ICES WEEE

*only available with Intel® Xeon® D 2187NT

Power consumption efficiency

Power consumption	x3 AI Inference Power Consumption 1.73 Tops/W vs. 0,67 T/W
-------------------	--

For more information: atos.net/BullSequanaEdge

Atos, the Atos logo, Atos Syntel, and Unify are registered trademarks of the Atos group. May 2019. © 2019 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.