

PRODUCT BRIEF

Lanner LEC-2290E
NVIDIA EGX AI Platform

Lanner

Reliable, Powerful Edge AI Computing for Accelerated real-time AI with NVIDIA® EGX™ Platform and Tesla T4 GPU

Lanner and NVIDIA create EGX platform for edge AI inference

Modern AI solutions are being built around existing highly flexible cloud architectures requiring compatible hardware and software that easily integrates into existing IT workflows

Artificial Intelligence is expanding and capabilities at the edge demand reliable, fast processing, and low latency to deliver key insights. Businesses seek to leverage powerful cloud-native approaches with all their solutions, while simplifying edge hardware deployments and orchestration. Emerging AI opportunities and workloads continue to drive the need for better tools and technology to stay ahead of competing solutions.

Lanner recognizes the challenges in these types of edge deployments, and has partnered with AI industry leader NVIDIA to deliver a solution built to handle modern AI inference from edge-to-cloud. Based on Nvidia's EGX™ platform, the LEC-2290E is built for diverse edge AI workloads, coming pre-configured with the efficient Tesla T4 tensor core GPU.


Reliable, scalable edge AI inference is at the heart of many real-world use cases. Increasingly intelligent manufacturing, retail outlets and smart cities all harness edge intelligence for automation, analytics, detection, user experience, demand prediction and much more.

Lanner

NVIDIA T4 GPU Performance

Single-Precision:	8.1 TFLOPS
Mixed-Precision: (FP16/FP32)	65 TFLOPS
INT8:	130 TOPS
INT4:	260 TOPS

Raw throughput is just one of seven factors to consider in the overall evaluation of a system. To get a complete picture of performance you must take into account all of the following:



Secure and powerful edge AI

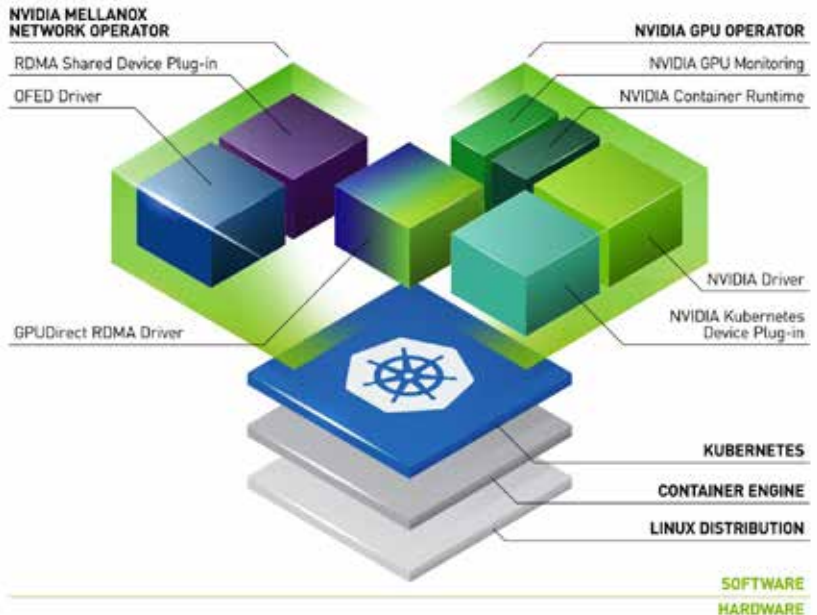
The NVIDIA EGX AI platform brings an easy-to-deploy, cloud-native software stack and management service together with GPU accelerated Lanner hardware

Businesses deploy flexible edge solutions that deliver the performance, reliability, hardware security and ease-of-use that fall in-line with their current investments in IT and cloud infrastructure.

With the NVIDIA EGX stack, traditional IT issues of installing software on remote systems disappear. Using Helm charts, containers, and continuous integration and continuous delivery (CI/CD), organizations can deploy updated AI containers effortlessly in minutes.

Lanner's LEC-2290 was built in collaboration with NVIDIA to create a robust edge AI inference node the power and simplicity of NGC-ready tests, EGX stack, drivers and hardware to fast-track new edge AI deployments.

The EGX stack architecture is supported by leading hybrid-cloud platform partners: Canonical, Cisco Container Platform (CCP), Microsoft Azure, Nutanix, Red Hat, and VMware.



LEC 2290E

Typical Work-flow of GPU powered AI deployments from edge to data center



Key product benefits

- Software-defined AI hardware
- NGC-Ready Edge Appliance*
- Cloud native software stack & orchestration
- Highly scalable
- Built-in security
- Performance optimized

*NGC-Ready Qualification tests are in-progress

Lanner 2290E Targeted Application Segments: Retail, Manufacturing, Smart Cities

From 2020-2025 AI is expected to grow by 25 billion in retail, manufacturing, and smart cities, with the demand for better, faster AI edge processing at its highest point. Below are examples of how NVIDIA EGX can be utilized within these industries to bring cost-effective optimizations and valuable insights for today.



Intelligent Retail

With AI-enabled intelligent stores, retailers are reducing shrinkage, eliminating stockout, and gaining visibility into in-store customer behavior to optimize merchandising. These stores leverage data from cameras and sensors to provide valuable analytics that enable smart decision-making, improve operations, and increase efficiency. Additionally, the same infrastructure can be used for a faster customer checkout experience, including fully automated checkout systems.

Application Segments: In-store analytics, autonomous shopping, asset protection, customer experience

Efficient Manufacturing

The flexibility of automation and fast-paced technologies like 3D printing and ecommerce have greatly increased capabilities & diversity, along with overall complexity for manufacturers. To meet these demands manufacturers are relying on powerful AI to increase efficiency and improve logistics.

Cutting edge AI implementations have enabled the some of the largest manufacturers to manage the immense influx of millions of parts around the world every day. Machine vision driven robotics have optimized manufacturing of increasingly intricate and personalized products.



Application Segments: Manufacturing & logistics

Smart Cities

Populous cities are deploying powerful autonomous AI solutions to reduce expenditure and increase revenue streams.

Problems like vehicle-tax collection on such scales require machine vision AI solutions that can handle traffic monitoring and vehicle/license plate recognition. Rising public safety concerns with the ongoing pandemic have accelerated the creation and real-world deployment of crowd and safety equipment detection solutions.

Application Segments: Traffic control, crowd detection





■ Features

- Intel® Xeon™ E-2278GEL 8C/16T
- NVIDIA T4 Tensor Core GPU for AI acceleration
- Pre-installed 64GB Wide temp SODIMM, 128Gb mSATA and 2.5" 256Gb SATA
- TPM 2.0 and IPMI modules
- 2x RJ45 GbE LAN, 4x PoE, 4x USB3.0, 6x COM Ports, 8x DI & 8x DO
- 1x PCIe *16, 1x Mini-PCIe (PCIe + USB2.0) w/ Nano-SIM, 1x B Key M.2 (PCIe + USB3.0) w/ Nano-SIM

■ Specifications

System Processor

CPU	Socketed Intel® Xeon™ E-2278GEL
Frequency	2.0GHz, 3.9 GHz Turbo
Core Number	8C/16T
BIOS	AMI 256Mbit SPI Flash
Chipset	BIOS FH82C246
Memory	
Technology	DDR4 2133/2400 SO-DIMM
Capacity	2x32GB Wide temp (64GB Max)
Socket	2x 260-pin SO-DIMM
Graphic	
Controller	Intel® UHD Graphics 630
HDMI	2
Audio	
Codec	TSI 92HD73C
Interface	1x Mic-in, 1x Line-out
Ethernet	
Controller	Intel i210IT
Speed	10/100/1000 Mbps
Interface	2x GbE RJ45, 4x PoE RJ45 @IEEE 802.3at (Total PoE Budget of 60W)
Storage	
Type(s)	SATA III
Expansion	2x Removable HDD/SSD External Slot w/ RAID
I/O	
Serial Port	6x D-Sub RS232/422/485
Digital I/O	8x DI, 8x DO (12V@100mA)
USB 2.0	-
USB 3.0	4x Type A
Power-On/Reset Button	1x Power-On, 1x Reset Button
Remote Power Switch	1x 2-pin Remote Power Switch
LED	Power/Storage/LTE/Wifi
Mechanical	
Dimension (W x H x D)	275 x 115 x 225 mm
Construction	Aluminum Extrusion & SGCC
Mounting	Wallmount



Tesla T4 Tensor Core GPU

GPU Architecture	NVIDIA Turing
NVIDIA Turing Tensor	320
NVIDIA CUDA® Cores	2,560
GPU Memory	16 GB GDDR6, 300 GB/sec
ECC	Yes
Interconnect Bandwidth	32 GB/sec
System Interface	x16 PCIe Gen3
Thermal Solution	Passive
Power	70W
Compute APIs	CUDA, NVIDIA TensorRT™, ONNX
Expansion Interface	
Mini-PCIe	1x Mini-PCIe (PCIe + USB2.0) w/ Nano-SIM
PCIe	1x PCIe *16 (In Use)
M.2	1x B Key M.2 w/ Nano-SIM
Misc	
Watchdog Timer	Yes
TPM	TPM 2.0
IPMI	Yes
Power	
Power Type	ATX
Power Supply Voltage	9~30 Vdc
Connector	1x 4-pin Terminal Block
Environment	
Operating Temperature	0°C~50°C (see *1)
Storage Temperature	-40°C~70°C
Relative Humidity	10%~90% (Non-condensing)
Vibration	IEC 60068-2-64, 0.5Grms, random 5 ~500 Hz, 40 mins/axis

*1 In compliance with NVIDIA® Tesla T4 warranty policy, an operating temperature of 0°C~50°C is required for systems with Tesla T4 installed.

Built-in hardware security and manageability

Hardware security is a major concern in edge deployment scenarios, with hardware-enhanced security being the gold-standard for professional-grade edge solutions. The LEC-2290E uses Intel's Coffee-Lake Xeon CPU, bringing numerous security enhancing and accelerating features. On-board TPM 2.0 for security and IPMI module for Manageability.



Intelligent Platform Management Interface (IPMI): IPMI offers system administrators to manage and monitor computing platforms with standardized interface and protocol. Due to its messaging and hardware-based nature, IPMI works independently from the operating system, so that system administrators are able to remotely manage and monitor computing platform status.



Trusted Platform Module (TPM): The TPM is a secure crypto-processor that assist with functions such as generating, storing and limiting the use of cryptographic keys. Keys generated by this crypto-processor provide tamper resistance at the hardware-level.

Performance and power efficiency at the very edge

NVIDIA T4 enterprise GPUs supercharge the world's most trusted mainstream servers. Its low-profile, 70-watt (W) design is powered by NVIDIA Turing™ Tensor Cores, delivering revolutionary multi-precision performance to accelerate a wide range of modern applications, including machine learning, deep learning, and virtual desktops.



Small-form-factor, 70-watt (W) GPU design: NVIDIA T4 optimized for scale-out servers, providing an incredible 50X higher energy efficiency compared to CPUs, drastically reducing operational costs. In the last two years, NVIDIA's inference platform has increased efficiency by over 10X, and it remains the most energy-efficient solution for distributed AI training and inference.



Data center class edge GPU: The T4 is an ideal universal accelerator for distributed edge computing environments. Revolutionary multi-precision performance accelerates deep learning and machine learning training and inference, video transcoding, and virtual desktops. NVIDIA T4 supports all AI frameworks and network types, delivering dramatic performance and efficiency that maximize the utility of at-scale deployments.

Lanner

About Lanner

Founded in 1986, Lanner is an ISO 9001 accredited organization with headquarters in Taipei, Taiwan, and offices in USA, Canada, and China. With over 30 years of experience in system and board hardware engineering, Lanner provides high-performance, reliable, and cost-effective computing platforms. Lanner is most renowned for its range of Intel architecture-based and RISC network appliances.

Lanner Electronics Inc. (TAIEX 6245) provides design, engineering, and manufacturing services for advanced network appliances and rugged applied computing platforms for system integrators, service providers, and application developers.

More information available at www.lanner-america.com.