FOR IMMEDIATE RELEASE

EMASS ECS-DoT-22 Delivers Up to 56× Lower Energy per Inference than Closest Published Peers on MLPerf® Tiny

Accelerator-native models beat reference accuracy on all four MLPerf Tiny workloads, run up to 9× faster, and bring keyword spotting and anomaly detection below one microjoule

LOS ANGELES (Nov. 12, 2026) – EMASS, a Nanoveu subsidiary specializing in ultra-efficient edge AI semiconductor technology, today released new benchmark results for its ECS-DoT-22 edge AI system-on-chip (SoC) across all four MLPerf Tiny workloads: keyword spotting, visual wake words, image classification, and anomaly detection. Running new models designed for the ECS-DoT deep-learning accelerator (DLA), the chip beat the MLPerf Tiny reference accuracy on every official test set, ran up to 9× faster than the reference models on the same silicon at 50 MHz, and posted the lowest published energy per inference on every workload.

Against the closest published peer on each workload, ECS-DoT-22's new models use 17× to 56× less energy per inference, run 1.2× to 11× faster, and cut energy-delay product (EDP, energy × latency) by 20× to 560×. Measured energy per inference is 0.56 µJ for keyword spotting and 0.39 µJ for anomaly detection. Even on the unmodified reference models, ECS-DoT-22 uses 10× to 19× less energy than the Syntiant Core 2.

Workload	Accuracy (new / reference)	Latency (new / reference)	Energy per inference	Energy / EDP gain vs closest peer
Keyword spotting	92.13% / 91.66%	434 µs / 3.9 ms	0.56 µJ	56× / 560× (Syntiant Core 2)
Visual wake words	86.30% / 84.80%	3.47 ms / 4.2 ms	3.94 µJ	18× / 67× (Syntiant Core 2)
Image classification	88.25% / 86.97%	1.39 ms / 6.2 ms	2.67 µJ	38× / 440× (Syntiant Core 2)
Anomaly detection	AUC 0.861 / 0.8595	397 µs / 1.2 ms	0.39 µJ	17× / 20× (BrainChip AKD1000)

Measured on ECS-DoT-22 at 50 MHz and 0.66 V core supply on the evaluation board, over the full official MLPerf Tiny test sets; company-measured, unverified MLPerf Tiny results. Closest peer is the lowest published energy per inference on each workload among vendors in current production, at that peer's lowest-energy operating point: Syntiant Core 2 (MLCommons-verified, 0.9 V / 30.7 MHz) and BrainChip AKD1000 (unverified). At Syntiant's faster 1.1 V / 98.7 MHz point, EDP gains are 270× (keyword spotting), 29× (visual wake words), and 190× (image classification).

For product designers, the results translate directly into battery life and user experience:

- **AI drops out of the power budget**
  At ten keyword-spotting inferences a second, the compute cost is about 6 µW, against about 315 µW at Syntiant Core 2's lowest-energy point, and less than a thousandth of what a Bluetooth or Wi-Fi radio uses to send the same audio to the cloud. Always-on listening or watching stops being the line item that sizes the battery.
- **The chip finishes sooner, then sleeps**
  Sub-millisecond keyword spotting and anomaly detection leave the accelerator idle almost all the time, with headroom to run several models on one sensor frame.

- **Fewer false triggers**
  Higher-than-reference accuracy on every workload means fewer missed wake words and fewer false person detections.

EMASS reports two sets of numbers: the unmodified MLPerf Tiny reference models, for a like-for-like comparison with closed-division results, and accelerator-native models held to at least the reference accuracy on the same official test sets. No product team ships the reference networks, which were written for general-purpose microcontrollers, and vendors publishing outside MLCommons, including BrainChip and Analog Devices, already report results on networks shaped for their own accelerators. The accelerator-native model is what a customer will receive from the EMASS IDE. On the reference models alone, ECS-DoT-22 already uses 10× to 19× less energy than the Syntiant Core 2; the accelerator-native models widen that lead rather than create it.

“Until now every ECS-DoT benchmark used networks written for microcontrollers. These models were built for our accelerator, and the difference shows: more accurate than the reference on every workload, up to nine times faster, and keyword spotting and anomaly detection below a microjoule,” said Mohamed M. Sabry Aly, Founder and Chief Technology Officer at EMASS. “This is the class-leading performance the next generation of always-on devices needs, across hearables, wearables, smart devices, industrial sensing, robotics and drones, and with the evaluation kits now on DigiKey, it is performance customers can put on their own bench.”

The results are built to be reproduced. The ECS-DoT-22 evaluation board (ECS-DoT-22-EVB) is available to order now from DigiKey, and the four benchmark models, ready-to-open IDE projects, and the host scripts behind every table will be released with the EMASS IDE and AI compiler later this month, so customers and independent third parties can repeat the tests.

Notes

1. ECS-DoT-22 results are company-measured and have not been verified by MLCommons Association or been through its MLPerf® review process. Reference-model results follow a methodology based on the MLPerf Tiny Closed Division; accelerator-native results use modified models consistent with an Open Division-style approach. Neither set is an official MLCommons-verified submission.
2. Unverified results may use measurement methodologies, model architectures, or workload implementations that differ from those prescribed for verified MLPerf® results, so comparisons between verified and unverified results may not be directly equivalent.
3. Energy per inference covers the processor core, on-chip memory, and deep-learning accelerator over one inference, measured on the ECS-DoT-22 evaluation board. It excludes peripherals and I/O interfaces and does not represent total device or system-level energy.
4. Peer results verified by MLCommons Association: MLPerf® Tiny v1.0 Closed (Silicon Labs), v1.0 to v1.2 Closed (Syntiant), v1.2 Closed (STMicroelectronics), and v0.7 Closed (Renesas; Silicon Labs anomaly detection), retrieved from mlcommons.org on Nov. 5, 2026. BrainChip AKD1000 results are unverified, from BrainChip’s January 2023 white paper. Comparisons use identified published results from different benchmark rounds and do not necessarily represent the latest products available from each manufacturer.
5. The MLPerf name and logo are trademarks of MLCommons Association in the United States and other countries. All rights reserved. Unauthorized use strictly prohibited.

###

FOR IMMEDIATE RELEASE

About EMASS

EMASS – a subsidiary of Nanoveu Ltd (ASX: NVU) – is an advanced semiconductor company specializing in ultra-low-power AI system-on-chip (SoC) solutions for edge computing. The company's flagship ECS-DoT chip delivers high-performance AI processing for vision, audio, and sensor data directly on-device, maximizing energy efficiency through its RISC-V architecture and non-volatile memory technologies. This always-on intelligence solution is optimized for power- and space-constrained applications including drones, wearables, healthcare devices and industrial IoT systems. For more information, visit nanoveu.com/emass.

About Nanoveu

Nanoveu is a listed company advancing human–machine experiences at the edge through a portfolio that spans ultra-low-power AI and glasses-free 3D technologies. Its subsidiary EMASS designs advanced system-on-chip (SoC) solutions that deliver efficient, scalable on-device AI for smart devices, IoT applications and 3D content transformation – enhancing Nanoveu's reach across rapidly growing AI, edge computing and 3D content markets. EyeFly3D™ is Nanoveu's end-to-end platform for glasses-free 3D, uniting proprietary screen technology with sophisticated content processing software and, now, EMASS's ultra-low-power SoC to bring immersive 3D to a wide range of devices and industries. The Company also develops and markets an advanced range of self-disinfecting and hydrophobic films and coatings under the Nanoshield™ brand, designed for applications including large-scale CSP and photovoltaic solar installations. Together, Nanoveu's businesses deliver practical innovation that makes devices smarter, environments safer and experiences more immersive.

Media Contact
Aaron Chong
Business Analyst, EMASS
aaron@nanoveu.com